

Juniper Networks JUNOS® 9.6 Software Release Notes

Release 9.6R4
01 June 2010
Revision 4

These release notes accompany Release 9.6R4 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 Ethernet 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 Juniper Networks M Series Multiservice Edge Routers, MX Series Ethernet Service Routers, and T Series Core Routers	6
New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers	6
Class of Service	6
High Availability	7
Interfaces and Chassis	9
JUNOS XML API and Scripting	15
Layer 2 Ethernet Services	20
Multicast	21
Network Management	25
Platform and Infrastructure	26
Routing Policy and Firewall Filters	27
Routing Protocols	30
Services Applications	31
Subscriber Access Management	37
User Interface and Configuration	43

VPNs	45
Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers	46
Forwarding and Sampling	46
Interfaces and Chassis	46
MPLS Applications	48
Platform and Infrastructure	48
Routing Protocols	48
Services Applications	50
Subscriber Access Management	51
User Interface and Configuration	54
VPNs	54
Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers	54
Current Software Release	55
Previous Releases	76
Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers	96
Changes to the JUNOS Documentation Set	96
Errata	96
Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers	103
Basic Procedure for Upgrading to Release 9.6	104
Upgrading a Router with Redundant Routing Engines	106
Upgrading the Software for a Routing Matrix	106
Upgrading Using ISSU	108
Upgrading from JUNOS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR	108
Downgrade from Release 9.6	109
JUNOS Software Release Notes for Juniper Networks SRX Series Services Gateways	110
New Features in JUNOS Release 9.6 for SRX Series Services Gateways	110
Software Features	111
Hardware Features—SRX100 Services Gateway	119
Hardware Features—SRX210 Services Gateways	122
Hardware Features—SRX5600 and SRX5800 Services Gateways	122
Changes In Default Behavior and Syntax in JUNOS Release 9.6 for SRX Series Services Gateways	123
Chassis Cluster	123
CLI	123
DHCP	123
Flow and Processing	123
Interfaces and Routing	124
Intrusion Detection and Prevention (IDP)	124
J-Web	124
Management and Administration	125

Known Limitations in JUNOS Release 9.6 for SRX Series Services	
Gateways	125
[accounting-options] Hierarchy	125
Chassis Cluster	125
Command-Line Interface (CLI)	125
Flow and Processing	126
Hardware	126
IGMP	127
Interfaces and Routing	127
Intrusion Detection and Prevention (IDP)	129
NetScreen-Remote	130
System	130
Issues in JUNOS Release 9.6 for SRX Series Services Gateways	130
Outstanding Issues In JUNOS Release 9.6 for SRX Series Services	
Gateways	130
Resolved Issues in JUNOS Release 9.6 for SRX Series Services	
Gateways	142
Errata and Changes in Documentation for JUNOS Release 9.6 for SRX	
Series Services Gateways	143
Application Layer Gateways (ALGs)	144
Attack Detection and Prevention	144
Chassis Cluster	144
CLI Reference	145
CompactFlash Card Support	146
DLSw	146
Flow	146
Incorrect Administration Features Support Information in	
Documentation	147
Incorrect Security Features Support Information in	
Documentation	147
Installing Software Packages	147
Intrusion Detection and Prevention (IDP)	149
J-Web	149
Network Management	149
Screens	151
JUNOS Software Release Notes for Juniper Networks J Series Services	
Routers	152
New Features in JUNOS Release 9.6 for J Series Services Routers	152
Software Features	153
Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series	
Services Routers	156
Configuration	157
Management and Administration	157
Security	157
Known Limitations in JUNOS Release 9.6 for J Series Services	
Routers	157
Chassis Cluster	158
IGMP	158
Interfaces and Routing	158
Intrusion Detection and Prevention (IDP)	159
J-Web	159

SNMP	159
Unified Threat Management (UTM)	159
Issues in JUNOS Release 9.6 for J Series Services Routers	159
Outstanding Issues In JUNOS Release 9.6 for J Series Services Routers	160
Resolved Issues in JUNOS Release 9.6 for J Series Services Routers	164
Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers	165
Application Layer Gateways (ALGs)	166
CLI Reference	166
DLSw	166
Flow	166
Intrusion Detection and Prevention (IDP)	166
J-Web	167
Screens	167
Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers	168
Transceiver Compatibility	168
Power and Heat Dissipation Requirements for J Series PIMs	168
Supported Third-Party Hardware for J Series Services Routers	168
J Series CompactFlash and Memory Requirements	169
Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers	170
JUNOS Software Release Notes for EX Series Switches	170
New Features in JUNOS Release 9.6 for EX Series Switches	170
Hardware	171
Access Control and Port Security	171
Ethernet Switching	171
Layer 3 Protocols	171
Management and RMON	172
Packet Filters	172
Virtual Chassis	172
Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches	173
Limitations in JUNOS Release 9.6 for EX Series Switches	173
Class of Service	173
Infrastructure	173
Interfaces	173
Outstanding Issues in JUNOS Release 9.6 for EX Series Switches	174
Bridging, VLANs, and Spanning Trees	174
Class of Service	175
Firewall Filters	175
Hardware	175
Infrastructure	175
Interfaces	177
Layer 2 and Layer 3 Protocols	178
Resolved Issues in JUNOS Release 9.6 for EX Series Switches	178
Access Control and Port Security	178
Infrastructure	178

Errata in Documentation for JUNOS Release 9.6 for EX Series	
Switches	179
Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series	
Switches	179
Upgrading or Downgrading from JUNOS Release 9.4R1 for EX Series	
Switches	179
Upgrading from JUNOS Release 9.3 to Release 9.6 for EX Series	
Switches	180
Upgrading from JUNOS Release 9.2 to Release 9.6 for EX Series	
Switches	180
Downgrading from JUNOS Release 9.6 to Release 9.2 for EX4200	
Switches	181
JUNOS Documentation and Release Notes	183
Documentation Feedback	183
Requesting Technical Support	183
Revision History	185

JUNOS Software Release Notes for Juniper Networks M Series Multiservice Edge Routers, MX Series Ethernet Service Routers, and T Series Core Routers

- New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 6
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 46
- Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 54
- Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers on page 96
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 103

New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers

The following features have been added to JUNOS Release 9.6. Following the description is the title of the manual or manuals to consult for further information.

Class of Service

- **IEEE 802.1p and DEI rewriting based on forwarding class and PLP at VPLS ingress PE (MX Series routers)**—Enables you to rewrite (mark) the IEEE 802.1p or DEI bits on frames at the VPLS ingress PE based on the value of the forwarding class and PLP established for the traffic on MX Series routers. You can rewrite either the outer tag only or both the outer and inner tag. When both tags are rewritten, both get the same value. To configure these rewrite rules, include the `ieee-802.1` statement at the `[edit class-of-service routing-instances name rewrite-rules]` hierarchy level. This feature is supported only on enhanced DPCs.
[Class of Service]
- **IEEE 802.1p and DEI rewriting based on forwarding class and PLP at VPLS ingress PE (IQ2 and IQ2E PICs)**—On routers with IQ2 or IQ2E PICs, enables you to rewrite (mark) the IEEE 802.1p or DEI bits on frames at the VPLS ingress PE based on the value of the forwarding class and PLP established for the traffic. You can rewrite either the outer tag separately or both the outer and inner tag. When both tags are rewritten, both get the same value.
To configure these rewrite rules, include the `ieee-802.1` statement at the `[edit class-of-service routing-instances name rewrite-rules]` hierarchy level.
[Class of Service]
- **BA classification based on the inner VLAN tag (IQ2 and IQ2E PICs)**—On routers with IQ2 or IQ2E PICs, enables you to perform behavior aggregate (BA) classification based on the inner VLAN tag. To configure BA classification, include the `inner` option at the `[edit class-of-service interfaces interface-name unit logical-unit-number classifiers ieee-802.1 classifier-name vlan-tag]` hierarchy level.
[Class of Service]

- **User-defined DSCP classification (queuing PFEs only)**—On MX, M120, and M320 routers with Enhanced Type III FPCs only, you can configure user-defined DSCP-based BA classification for MPLS interfaces (this feature is not available for IQE PICs) or VPLS/L3VPN routing instances (LSI interfaces). The DSCP-based classification for MPLS packets for Layer 2 VPNs is not supported. To classify MPLS packets on the routing instance at the egress PE, include the `dscp` or `dscp-ipv6` statements at the `[edit class-of-service routing-instances routing-instance-name classifiers]` hierarchy level. To classify MPLS packets at the core-facing interface, apply the classifier at the `edit class-of-service interface interface-name unit unit-name classifiers (dscp | dscp-ipv6) classifier-name family mpls` hierarchy level (this feature is not available for IQE PICs or on MX Series routers when ingress queuing is used).

[Class of Service]

- **Forwarding class map for unicast and multicast traffic and a user-configured queue number for egress interfaces**—For the IQ, IQ2, IQE, LSQ, and ATM2 PICs in the T320, T640, and T1600 routers, you can configure a forwarding class map for unicast and multicast traffic and a user-configured queue number for an egress interface. DSCP rewrites are not allowed on these interfaces. You cannot apply the forwarding class map on an LSQ physical interface because per-unit scheduling is enabled by default and cannot be disabled manually. The interface-specific forwarding-class map is configured at the `[edit class-of-service forwarding-classes-interface-specific forwarding-class-map-name class class-name queue-num queue-number (restricted-queue queue-number)]` hierarchy level. You apply the forwarding-class map to the logical unit level at the `[edit interfaces interface-name unit logical-unit-number output-forwarding-class-map forwarding-class-map-name]` hierarchy level. For PICs that are restricted to four queues, you can control the queue assignment with the `restricted-queue` option or rely on the routing platform to override the queue assignment using modular arithmetic.

[Class of Service, Multicast]

High Availability

- **New configuration statement to control VRRP advertisements**—A new configuration statement `vrrp-inherit-from` at the `[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-number]` hierarchy level enables you to control VRRP advertisements from various VRRP groups configured on different subnets of a VLAN. When the `vrrp-inherit-from` configuration statement is included in the configuration, only the master VRRP group from which the other master VRRP groups are inheriting the state sends out VRRP advertisements. The groups inheriting the state do not send any VRRP advertisements as the state is maintained only on the group from which the state is inherited. If the `vrrp-inherit-from` statement is not configured, each of the VRRP master groups in the various subnets on the VLAN sends out separate VRRP advertisements and adds to the traffic on the VLAN.

When you have included the `vrrp-inherit-from` statement for a VRRP group, the VRRP group inherits the following parameters from the specified group:

- `advertise-interval`
- `authentication-key`

- authentication-type
- fast-interval
- preempt | no-preempt
- priority
- track interfaces
- track routes

However, you can configure the **accept-data | no-accept-data** statement for the group to specify whether the interface should accept packets destined for the virtual IP address.

[High Availability]

- **Nonstop active routing support for LDP OAM features**—Extends nonstop active routing support for LDP Operation, Administration, and Maintenance (OAM) features. Nonstop active routing support for LDP OAM features ensures that both egress and ingress LSPs maintain OAM state information, including information about BFD sessions during and after Routing Engine switchover, and minimizes control plane churn and traffic loss during and after the Routing Engine switchover.

[High Availability]

- **Graceful Routing Engine switchover (GRES) support for TX Plus routing matrix**—JUNOS Release 9.6 extends GRES support to the TX Matrix Plus router.



NOTE: TX Matrix Plus routers and T1600 routers that are configured as part of a routing matrix do not currently support nonstop active routing.

[High Availability]

- **Graceful Routing Engine switchover (GRES) and nonstop active routing support for pseudowire configurations**—JUNOS Release 9.6 extends GRES and nonstop active routing support to Layer 2 circuit and LDP-based VPLS pseudowire configurations.

[High Availability]

- **Distributed periodic packet management (PPMD) support for VRRP**—Typically, VRRP advertisements are sent by the VRRP process (vrrpd) on the master VRRP router at regular intervals to let other members of the group know that the VRRP master router is operational.

When the VRRP process is busy and does not send VRRP advertisements, the backup VRRP routers might assume that the master router is down and take over as the master router, causing unnecessary flaps as the takeover occurs, irrespective of the fact that the original master is still active and available, and might restart sending advertisements after the load has decreased. To address this problem and to reduce the load on VRRPD, the JUNOS Software now uses the Periodic Packet Management process (PPMD) to send VRRP advertisements on behalf of VRRPD.

However, you can further delegate the job of sending VRRP advertisements to the distributed PPMD that resides on the Packet Forwarding Engine. The ability to delegate the sending of VRRP advertisements to the distributed PPMD ensures that the VRRP advertisements are sent even when PPMD—which is now responsible for generating VRRP advertisements—is busy, and prevents the possibility of false alarms when PPMD is busy.

To configure PPMD to send VRRP advertisements when VRRPD is busy, include the `set delegate-processing` statement at the `[edit protocols vrrp]` hierarchy level.

```
[edit protocols vrrp]
  set delegate-processing;
```

[*High Availability*]

- **Unified ISSU support on additional hardware**—JUNOS Release 9.6 extends the unified ISSU support to the following PICs:
 - 4-port Ethernet (PD-4XGE-XFP) PIC (only on T640 and T1600 routers)
 - 1-port Channelized OC48/STM16 Enhanced IQ2 PIC (PB-1CHOC48-STM16-IQE)

[*High Availability*]

Interfaces and Chassis

- **IPv6 support for unnumbered Ethernet interfaces (M Series and MX Series routers)**—Enables you to configure IPv6 processing on an Ethernet interface without assigning an explicit IPv6 address to the interface. Using unnumbered interfaces allows a single subnet to be shared across multiple interfaces. This feature also supports nonsubnetted Ethernet for IPv6 (/128 addressing).

When you configure an unnumbered Ethernet interface, you configure the interface to borrow an address from a donor interface. To specify the name of the donor interface, include the `unnumbered-address interface-name` statement at the `[edit interfaces interface-name unit logical-unit-number family inet6]` hierarchy level. The `unnumbered-address interface-name` statement is also supported at the `[edit dynamic-profiles profile-name interfaces interface-name]` and `[edit logical-systems logical-system-name interfaces interface-name]` hierarchy levels.

The output of the `show interface` operational command now includes information for IPv6 unnumbered Ethernet interfaces.

[*Network Interfaces, Routing Protocols, Interfaces Command Reference, Subscriber Access*]

- **802.1ag optional type, length, and value (TLV) support**—Provides configuration support for the Port Status TLV and the Interface Status TLV on M120, M320, and MX Series routers. Configuring the Port Status TLV allows the operator to control the transmission of the Port Status TLV in connectivity fault management PDUs. The Interface Status TLV indicates the status of the interface on which the MEP transmitting the CCM is configured (which is not necessarily the interface on which it resides), or the next-lower interface in the IETF RFC 2863 IF-MIB.

To configure the Port Status TLV, include the `port-status-tlv` option at the `[edit protocols oam ethernet connectivity-fault-management maintenance-domain identifier`

maintenance-association *identifier* continuity-check] hierarchy level. To configure the Interface Status TLV, include the `interface-status-tlv` option at the [edit protocols oam ethernet connectivity-fault-management maintenance-domain *identifier* maintenance-association *identifier* continuity-check] hierarchy level.

You can use show commands to display the different values of the Port Status TLV, Interface Status TLV, and action profile information. To display information about a local MEP, use the `show oam ethernet connectivity-fault-management mep-database maintenance-domain identifier maintenance-association identifier local-mep identifier remote-mep identifier` command.

To display information about a remote MEP, use the `show oam ethernet connectivity-fault-management mep-database maintenance domain identifier maintenance-association identifier remote-mep identifier` command.

To display information about an event configured under an action profile that occurs for a remote MEP, including the action taken and the corresponding time, you can display the MEP database output for that remote MEP using the `show oam ethernet connectivity-fault-management mep-database maintenance domain identifier maintenance-association identifier remote-mep identifier` command.

When all configured events (under an action profile) are cleared, then the action taken gets reverted (for example, the down interface is up) and corresponding time can be displayed using the `show oam ethernet connectivity-fault-management mep-database maintenance domain identifier maintenance-association identifier remote-mep identifier` command.

[Network Interfaces]

- **Channelized ATM support on 12-port T1/E1 CE and 4-port Channelized OC3/STM1 circuit emulation interfaces**—On M7i, M10i, M120, M320, and M40e routers, support is provided for ATM pseudowires (RFC 4717) and packet encapsulations (RFC 2684).

CE PICs support the following ATM protocols:

- ATM over PWE3 (RFC 4717)
- ATM PWE3 via dynamic labels (LDP, RSVP-TE)

CE PICs provide ATM pseudowire encapsulation support for cell relay pseudowire or AAL5 pseudowire. Cell relay mode supports logical or physical interfaces. To configure a cell-relay pseudowire, use the `encapsulation atm-ccc-cell-relay` and `atm-l2circuit-mode cell` statements at the [edit interfaces at-fpc/pic/port unit n] hierarchy level. To configure an AAL5 pseudowire, use the `encapsulation atm-ccc-vc-mux` and `atm-l2circuit-mode aal5` statements at the [edit interfaces at-fpc/pic/port unit n] hierarchy level.



NOTE: The `encapsulation atm-ccc-cell-relay` command can be set at either the physical interface or logical interface. `atm-ccc-vc-mux` can only be set at the logical interface.

ATM OAM on CE interfaces supports the following OAM-FM cell types:

- F4 AIS (end-to-end)
- F4 RDI (end-to-end)

- F4 loopback (end-to-end)
- F5 loopback
- F5 AIS
- F5 RDI

ATM encapsulations provide congestion control via EPD thresholds on a per logical interface basis. For CE PICs, the EPD number specifies the number of packets (or frames, or cell bundles). To configure the EPD threshold, use the `epd-threshold <packets>` statement and `plp1 <packets>` option at the `[edit interfaces at-fpc/pic/port unit unit-number]` hierarchy level.

When you configure CE PICs for ATM, you must specify `atm-ce` as the PIC type in the `atm-options`. To configure the PIC type, use the `pic-type <atm-ce>` statement at the `[edit interfaces at-fpc/pic/port]` hierarchy level.

With CE PICs, ATM is also supported on E1 and T1 media. You can configure the E1 options or T1 options using the `e1-options` or `t1-options` statement at the `[edit interfaces at-fpc/pic/port]` hierarchy level.

[Network Interfaces]

- **Chassis operational mode commands supported on the TX Matrix Plus router**—The following CLI operational mode commands support chassis and fabric management operations on a TX Matrix Plus router:

Show commands:

- `show chassis alarms`
- `show chassis cip` (new command introduced for the TX Matrix Plus router)
- `show chassis craft-interface`
- `show chassis environment`
- `show chassis environment cb`
- `show chassis environment cip`
- `show chassis environment pem`
- `show chassis environment routing-engine`
- `show chassis environment scg`
- `show chassis environment sib`
- `show chassis ethernet-switch`
- `show chassis fabric plane`
- `show chassis fabric plane-location`
- `show chassis fabric topology`
- `show chassis firmware`

- show chassis hardware
- show chassis location
- show chassis routing-engines
- show chassis sibs
- show chassis spmb
- show chassis temperature-thresholds

Request commands:

- request chassis cb
- request chassis cip (new command introduced for the TX Matrix Plus router)
- request chassis fpm
- request chassis sib
- request chassis spmb

Clear command:

- clear chassis display message

Set command:

- set chassis display

Restart command:

- restart chassis-control

[System Basics and Services Command Reference]

- **New management Ethernet interface em0 for the TX Matrix Plus router and T1600 routers in a routing matrix**—The management Ethernet interface used for the TX Matrix Plus router and the T1600 routers in a routing matrix is **em0**. This interface provides an out-of-band method for connecting to the routers in the routing matrix.

**NOTE:**

- The Routing Engines in the TX Matrix Plus router and the T1600 routers (in the routing matrix) support only the management Ethernet interface, **em0**. They do not support the management Ethernet interface, **fxp0**.
- Automated scripts developed for standalone T1600 routers that are not in a routing matrix cannot be directly used on T1600 routers in a routing matrix. They might contain references to the **fxp0** management Ethernet interface. Before reusing the scripts on T1600 routers in a routing matrix, ensure that all commands in the scripts that reference the **fxp0** management Ethernet interface are updated to reference the **em0** management Ethernet interface.

[TX Matrix Plus Hardware Guide, Network Interfaces, System Basics, Interfaces Command Reference, System Basics and Services Command Reference, Software Installation and Upgrade Guide, High Availability]

- **Redundant Service PIC (RSP) hot standby (M Series and T Series platforms with MS-PICs and MX Series routers with MS-DPCs)**—Enables the use of a redundant pair of service PICs to operate in hot standby mode, facilitating faster switching to the standby PIC in a failover situation. Configuration changes are applied to both of the paired PICs simultaneously. This differs from the previously implemented warm standby functionality, in which the standby PIC receives configuration information at the time of failover. The maximum switchover time is 5 seconds. A typical switchover time would depend on the failure conditions and could be much less than the maximum.

To configure hot standby for a pair of service PICs, include the **hot-standby** statement at the **[edit interfaces rsp0 redundancy-options]** hierarchy level.

[Multiplay Solutions, Services Interfaces]

- **Aggregated interfaces support hierarchical queuing and shaping**—M120, MX Series, and T Series routers, with aggregated Ethernet IQ2 PICs in non-link-protect mode, support the following scheduler functions:
 - Per-unit scheduler
 - Hierarchical scheduler
 - Shaping at the physical and logical interface (aggregated interface) level

You can configure the **hierarchical-scheduler** mode on an aggregated Ethernet interface in non-link-protect mode using the **aggregated-ether-options** statement at the **[edit interfaces aeX]** hierarchy level. Prior to JUNOS Release 9.6, the hierarchical scheduler mode required the **aggregated-ether-options** statement **link-protection** option, otherwise a configuration error would result.

You can also specify the member link bandwidth derivation based on the equal division model (**scale**) or the replication model (**replicate**) using the **member-link-scheduler (scale | replicate)** option at the **[edit class-of-service interfaces aeX]** hierarchy level. The default setting is **scale**.

[Network Interfaces]

- **Support for the 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP) as a shared interface PIC on the JCS1200 platform**—On the JCS1200 platform, a single Physical Interface Card (PIC) can be used by different PSDs to forward and receive network traffic. With JUNOS Release 9.6, the 4-port 10-Gigabit Ethernet LAN/WAN PIC with XFP can be used as a shared interface.
[Protected System Domain]
- **802.1ag Ethernet OAM for CCC encapsulated packets**—On M120 and MX Series routers, CFM sessions on interfaces with CCC encapsulation can monitor L2VPN circuits for both L2 circuits and L2 VPNs.

CFM features supported on L2VPN circuits are as follows:

- Creation of UP/DOWN MEPs at any level on the customer edge (CE)-facing logical interface.
- Creation of MIPs at any level on the CE-facing logical interface.
- Support for continuity check, loopback, and linktrace protocol.
- Support for Y1731 Ethernet delay measurement protocol.
- Support for action profiles to bring the CE-facing logical interface down when loss of connectivity is detected.

To monitor an L2VPN circuit, CFM UP MEP can be configured on the CE-facing logical interfaces of the provider edge routers. To monitor the CE-PE attachment circuit, a CFM DOWN MEP can be configured on the customer logical interfaces of CEn-PEn routers.

To create a MIP on the CE-facing interface of the PE router, use the `interface (ge | xe)-fpc/pic/port.unit` statement at the `[protocols oam ethernet connectivity-fault-management maintenance-domain <identifier>]` hierarchy level.

[Network Interfaces]

- **Symmetrical load balancing on 802.3ad Link Aggregation (LAG) on MX Series routers**—MX Series routers with aggregated Ethernet PICs now support symmetrical load balancing on 802.3ad LAG. This feature is significant when two MX Series routers are connected transparently through Deep Packet Inspection (DPI) devices over a LAG bundle. DPI devices keep track of flows and require information of a given flow in both forward and reverse directions. Before symmetrical load balancing on 802.3ad LAG, the DPIs could misunderstand the flow, leading to traffic disruptions. Adding this feature means that a given flow of traffic (duplex) is ensured to hit the same devices in both directions.

This feature utilizes a mechanism of interchanging the source and destination addresses for a hash computation of fields, such as source address and destination address. The result of a hash computed on these fields is used to choose the link of the LAG. The hash computation for the forward and reverse flow must be identical. This is achieved by swapping source fields with destination fields for the reverse flow. The swapped operation is *complement hash computation* or *symmetric-hash complement* and the regular (or unswapped) operation is *symmetric-hash computation* or *symmetric-hash*. The swappable fields are: MAC address, IP address, and port.

You can now specify whether symmetric hash or complement hash is done for load-balancing traffic. To configure symmetric hash, use the `symmetric-hash` statement at the `[edit forwarding-options hash-key family inet]` hierarchy level. To configure symmetric-hash complement, use the `symmetric-hash <complement>` statement and option at the `[edit forwarding-options hash-key family inet]` hierarchy level.

These operations can also be performed at the PIC level by specifying a *hash key*. To configure a hash key at the PIC level, use the `symmetric-hash` or `symmetric-hash <complement>` statement at the `[edit chassis fpc fpc-slot pic pic-slot hash-key family inet]` and `[edit chassis fpc fpc-slot pic pic-slot hash-key family multiservice]` hierarchy levels.

The following restrictions apply:

- The Packet Forwarding Engine (PFE) can be configured to hash the traffic in either symmetric or complement mode. A single PFE complex cannot work simultaneously in both operational modes and such a configuration can yield undesirable results.
- The per-PFE setting overrides the chassis-wide setting only for the family configured. For the other families, the PFE complex still inherits the chassis-wide setting (when configured) or the default setting.
- Any change in the hash-key configuration at the PIC level requires a reboot of the FPC for the changes to take effect.
- This feature supports VPLS, INET, and bridged traffic only.
- This feature cannot work in tandem with the `per-flow-hash-seed <load-balancing>` option. It requires that all the PFE complexes configured in complementary fashion share the same seed. A change in the seed between two counterpart PFE complexes might yield undesired results.

In order for the symmetric or complement load balancing to work, the child member links of the LAGs should have identical order. To achieve this, configure the link-index while adding the child using the `[edit interfaces interface gigether-options 802.3ad link-index link-index-number]` command.

[*Network Interfaces, VPN, System Basics*]

JUNOS XML API and Scripting

- **Support for master source file for event script**—Enables you to refresh the local copy of an event script from a master copy stored in a central repository. To configure the location of the master source file for a specific event script, include the `source` statement at the `[edit event-options event-script file filename]` hierarchy level. Specify the source as a Hypertext Transfer Protocol (HTTP) URL, FTP URL, or secure copy (scp)-style remote file specification. You then can refresh the local event script from the master copy by including the `refresh` statement at either the `[edit event-options event-script file filename]` hierarchy level to refresh a single event script or at the `[edit event-options event-script]` hierarchy level to refresh all enabled event scripts from their configured source files.

In addition, you can refresh local event scripts from copies at an alternate location. Include the `refresh-from` statement and the source URL at the `[edit event-options event-script file filename]` hierarchy level to refresh a single event script or at the `[edit event-options event-script]` hierarchy level to refresh all enabled event scripts from copies at the specified source URL.

[*Configuration and Diagnostic Automation Guide*]

- **New JUNOS XML API operational request tag elements**—Table 1 on page 16 lists the JUNOS Extensible Markup Language (XML) operational request tag elements that are new in JUNOS Release 9.6, along with their corresponding CLI command and response tag element.

Table 1: New JUNOS XML Tag Elements and CLI Command Equivalents in JUNOS 9.6

Request Tag Element	CLI Command	Response Tag Element
<clear-bgp-damping>	clear bgp damping	NONE
<clear-bgp-neighbor>	clear bgp neighbor	NONE
<clear-bgp-table>	clear bgp table	NONE
<clear-cfm-delay-statistics>	clear oam ethernet connectivity-fault-management delay-statistics	NONE
<clear-cfm-linktrace-path-database>	clear oam ethernet connectivity-fault-management path-database	NONE
<clear-cfm-statistics>	clear oam ethernet connectivity-fault-management statistics	NONE
<clear-diameter-function>	clear diameter function	NONE
<clear-diameter-peer>	clear diameter peer	NONE
<clear-isis-adjacency-information>	clear isis adjacency	NONE
<clear-isis-database-information>	clear isis database	NONE
<clear-isis-overload-information>	clear isis overload	NONE
<clear-isis-statistics-information>	clear isis statistics	NONE
<clear-service-msp-flow-ipaction-table>	clear services flows ip-action	<service-msp-flow-drain-information>
<clear-system-services-reverse-information>	clear system services reverse	NONE
<get-bgp-rtf-information>	show bgp group rtf	<bgp-rtf-information>
<get-cfm-delay-statistics>	show oam ethernet connectivity-fault-management delay-statistics	<cfm-delay-statistics>
<get-cfm-forwarding-state-instance-information>	show oam ethernet connectivity-fault-management forwarding-state instance	<cfm-flood-instance-information>

Table 1: New JUNOS XML Tag Elements and CLI Command Equivalents in JUNOS 9.6 (continued)

Request Tag Element	CLI Command	Response Tag Element
<get-cfm-forwarding-state-interface-information>	show oam ethernet connectivity-fault-management forwarding-state interface	<cfm-flood-interface-information>
<get-cfm-interfaces-information>	show oam ethernet connectivity-fault-management interfaces	<cfm-interface>
<get-cfm-linktrace-path-database>	show oam ethernet connectivity-fault-management path-database	<cfm-linktrace-path-database>
<get-cfm-mep-database>	show oam ethernet connectivity-fault-management mep-database	<cfm-mep-database>
<get-cfm-mep-statistics>	show oam ethernet connectivity-fault-management mep-statistics	<cfm-mep-statistics>
<get-cfm-mip-information>	show oam ethernet connectivity-fault-management mip	<cfm-mip-information>
<get-cos-classifier-table-information>	show class-of-service forwarding-table classifier	<cos-classifier-table-information>
<get-cos-classifier-table-map-information>	show class-of-service forwarding-table classifier mapping	<cos-classifier-table-map-information>
<get-cos-forwarding-class-information>	show class-of-service forwarding-class	<cos-forwarding-class-information>
<get-cos-forwarding-class-map-interface-table-information>	show class-of-service forwarding-table forwarding-class-map mapping	<cos-forwarding-class-map-interface-table-information>
<get-cos-forwarding-class-map-table-information>	show class-of-service forwarding-table forwarding-class-map	<cos-forwarding-class-map-table-information>
<get-cos-fragmentation-map-information>	show class-of-service fragmentation-map	<cos-fragmentation-map-information>
<get-cos-fwtab-fabric-scheduler-map-information>	show class-of-service forwarding-table fabric scheduler-map	<cos-fwtab-fabric-scheduler-map-information>
<get-cos-interface-map-information>	show class-of-service interface	<cos-interface-information>
<get-cos-interface-set-map-information>	show class-of-service interface-set	<cos-interface-set-information>
<get-cos-l2tp-session-map-information>	show class-of-service l2tp-session	<cos-l2tp-session-information>
<get-cos-loss-priority-map-information>	show class-of-service loss-priority-map	<cos-loss-priority-map-information>
<get-cos-loss-priority-map-table-binding-information>	show class-of-service forwarding-table loss-priority-map mapping	<cos-loss-priority-map-table-binding-information>
<get-cos-loss-priority-map-table-information>	show class-of-service forwarding-table loss-priority-map	<cos-loss-priority-map-table-information>
<get-cos-multi-destination-information>	show class-of-service multi-destination	<cos-multi-destination-information>

Table 1: New JUNOS XML Tag Elements and CLI Command Equivalents in JUNOS 9.6 *(continued)*

Request Tag Element	CLI Command	Response Tag Element
<get-cos-policer-table-map-information>	show class-of-service forwarding-table policer	<cos-policer-table-map-information>
<get-cos-red-information>	show class-of-service forwarding-table drop-profile	<cos-red-information>
<get-cos-rewrite-information>	show class-of-service rewrite-rule	<cos-rewrite-information>
<get-cos-rewrite-table-information>	show class-of-service forwarding-table rewrite-rule	<cos-rewrite-table-information>
<get-cos-rewrite-table-map-information>	show class-of-service forwarding-table rewrite-rule mapping	<cos-rewrite-table-map-information>
<get-cos-routing-instance-map-information>	show class-of-service routing-instance	<cos-routing-instance-information>
<get-cos-scheduler-map-information>	show class-of-service scheduler-map	<cos-scheduler-map-information>
<get-cos-scheduler-map-table-information>	show class-of-service forwarding-table scheduler-map	<cos-scheduler-map-table-information>
<get-cos-shaper-table-map-information>	show class-of-service forwarding-table shaper	<cos-shaper-table-map-information>
<get-cos-table-information>	show class-of-service forwarding-table	<cos-table-information>
<get-cos-traffic-control-profile-information>	show class-of-service traffic-control-profile	<cos-traffic-control-profile-information>
<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-cos-virtual-channel-group-information>	show class-of-service virtual-channel-group	<cos-virtual-channel-group-information>
<get-cos-virtual-channel-information>	show class-of-service virtual-channel	<cos-virtual-channel-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-dhcp-relay-binding-information>	show dhcp relay binding	<dhcp-relay-binding-information>
<get-dhcp-relay-statistics-information>	show dhcp relay statistics	<dhcp-relay-statistics-information>
<get-dhcp-server-binding-information>	show dhcp server binding	<dhcp-server-binding-information>
<get-dhcp-server-statistics-information>	show dhcp server statistics	<dhcp-server-statistics-information>

Table 1: New JUNOS XML Tag Elements and CLI Command Equivalents in JUNOS 9.6 *(continued)*

Request Tag Element	CLI Command	Response Tag Element
<get-diameter-function-information>	show diameter function	<diameter-function-information>
<get-diameter-function-statistics>	show diameter function statistics	<diameter-function-statistics>
<get-diameter-information>	show diameter	<diameter-information>
<get-diameter-instance-information>	show diameter instance	<diameter-instance-information>
<get-diameter-network-element-information>	show diameter network-element	<diameter-network-element-information>
<get-diameter-network-element-map-information>	show diameter network-element map	<diameter-network-element-map-information>
<get-diameter-peer-information>	show diameter peer	<diameter-peer-information>
<get-diameter-peer-map-information>	show diameter peer map	<diameter-peer-map-information>
<get-diameter-peer-statistics>	show diameter peer statistics	<diameter-peer-statistics>
<get-diameter-route-information>	show diameter route	<diameter-route-information>
<get-fabric-queue-information>	show class-of-service fabric statistics	<fabric-queue-information>
<get-igmp-output-group-information>	show igmp output-group	<igmp-output-group-information>
<get-lldp-information>	show lldp	<lldp-information>
<get-lldp-information-detail>	show lldp detail	<lldp-information-detail>
<get-lldp-local-info>	show lldp local-information	<lldp-local-info>
<get-lldp-neighbors-information>	show lldp neighbors	<lldp-neighbors-information>
<get-lldp-remote-global-statistics>	show lldp remote-global-statistics	<lldp-remote-global-statistics>
<get-lldp-statistics-information>	show lldp statistics	<lldp-statistics-information>
<get-mld-output-group-information>	show mld output-group	<mld-output-group-information>
<get-multicast-pim-to-igmp-proxy-information>	show multicast pim-to-igmp-proxy	<multicast-pim-to-igmp-proxy-information>
<get-multicast-pim-to-mld-proxy-information>	show multicast pim-to-mld-proxy	<multicast-pim-to-mld-proxy-information>
<get-statistics-information>	show system statistics	<statistics>

Table 1: New JUNOS XML Tag Elements and CLI Command Equivalents in JUNOS 9.6 (*continued*)

Request Tag Element	CLI Command	Response Tag Element
<get-system-services-reverse-information>	show system services reverse	<system-services-reverse-information>
<reload-eedebg-action-profile>	request security datapath-debug action-profile reload-all	<message>
<request-monitor-ethernet-delay-measurement>	monitor ethernet delay-measurement	<ethdm-results>
<request-ping-ethernet>	ping ethernet	<ethping-results>
<request-traceroute-ethernet>	traceroute ethernet	<ethtraceroute-results>

[*JUNOS XML API Operational Reference*]

Layer 2 Ethernet Services

- **Support for Layer 2 services on routers with MS-DPCs (MX Series routers)**—Layer 2 link services are supported on MX Series routers with MS-DPCs containing MS-DPC PICs that eventually bundle PPP links from the Type 2 channelized SONET PICs.

To enable the Layer 2 service package for LSQ support, include the `service-package layer-2` statement at the `[edit chassis fpc slot-number pic pic-number adaptive-services]` hierarchy level.

[*System Basics, Services Interfaces*]

- **Link Layer Discovery Protocol (LLDP) support (MX Series routers only)**—Enables you to configure the LLDP protocol on interfaces on MX Series routers only. To configure and adjust default parameters, include the `lldp` statement at the `[edit protocols]` hierarchy level.

LLDP is disabled by default. Use the `enable` statement to enable LLDP and the `interfaces` statement to enable LLDP on all or some interfaces. Adjust the frequency that LLDP advertises with the `advertisement-interval` statement. The default is 30 seconds. Adjust the transmit delay that LLDP delays successive advertisements with the `transmit-delay` statement. The default is 2 seconds. Adjust the hold multiplier that LLDP uses to purge the cache or learned information with the `hold-multiplier` statement. The default is 4 (or 120 seconds with the default advertisement interval). Adjust the physical topology trap interval that LLDP sends SNMP traps about statistics with the `ptopo-configuration-trap-interval` statement. The default is 60 seconds. Adjust the physical topology maximum hold time that LLDP holds dynamic entries with the `ptopo-configuration-maximum-hold-time` statement. The default is 300 seconds.

[*Layer 2 Configuration Guide*]

- **Layer 2 address learning in logical systems (MX Series routers only)**—Enables Layer 2 address learning in logical systems for bridge domains and other

virtual-switch routing instances on MX Series routers only. To configure, include the `bridge-domains` or `switch-options` statements at the `[edit logical-systems logical-system-name]` hierarchy level.

[*Layer 2 Configuration Guide*]

Multicast

- **Support for multicast OIF mapping**—Multicast output interface (OIF) mapping enables you to configure multicast streams (data packets) so that they are directed to a different interface than multicast control packets (used for IGMP and MLD joins and leaves). If multiple control interfaces all request the same stream and are mapped to the same output interface, only one copy of the stream is transmitted.



NOTE: OIF mapping and reverse OIF mapping are not supported on the same customer or shared interface.

Configuring OIF mapping on a multicast interface requires the following general steps:

1. Specify an OIF map.

The OIF map is a routing policy statement that can contain multiple terms. To specify an OIF map, use the `policy-statement` statement at the `[edit policy-options policy-statement]` hierarchy level. The *from* clause in each term can select multicast flows based on the multicast group (using the `route-filter` keyword) and the source address (with the `source-address-filter` keyword). The *then* clause must accept or reject the term. However, the *then* clause is enhanced with a new `map-to-interface` statement at the `[edit policy-options policy-statement policy-name term term-name then actions]` hierarchy level. The `map-to-interface` statement sets a value similar to the existing metric or tag actions and requires you to specify either a logical interface (that is, any interface that multicast currently supports) or the keyword `self`. The `self` keyword specifies that multicast data packets are sent on the same interface as the control packets and no mapping occurs. If no term matches, then no multicast data packets are sent.

When creating OIF maps, keep the following in mind:

- If a physical interface is specified as the `map-to-interface` statement value (for example, `ge-0/0/0`), a ".0" is appended to the interface to create a logical interface (for example, `ge-0/0/0.0`).
- Routing policy is configured using the CLI per logical system. You cannot configure routing policies dynamically. As part of the policy, you must configure the `map-to-interface` statement statically.
- To use OIF mapping, you must also have IGMP, MLD, or PIM configured (see below for configuring IGMP or MLD as passive).
- You cannot map to a mapped interface.

- To use CAC with a mapped interface, you must configure it using the `maximum-bandwidth` statement (see the existing documentation for details about this statement).
 - Juniper Networks recommends that you configure policy statements for IGMP and MLD separately.
2. Associate the OIF map with an interface.

To associate a map with a logical interface, use the new `oif-map` statement at the `[edit protocols igmp interface interface-name]`, `[edit dynamic-profiles profile-name protocols igmp interface interface-name]`, `[edit protocols mld interface interface-name]`, or `[edit dynamic-profiles profile-name protocols mld interface interface-name]` hierarchy levels. The OIF map is applied to all IGMP or MLD requests received on the configured interface.



NOTE: If an OIF map is already configured on an interface, the new OIF map is added to the policy.

3. Disable QoS adjustment.

When configured, the OIF map performs QoS adjustment on the customer interface. QoS adjustment decreases the available bandwidth to the client interface by mapping multicast streams to the shared interface. This action occurs by default as long as a value is configured for the `maximum-bandwidth` statement at the `[edit routing-options multicast interface interface-name]` hierarchy level.

If you do not want to reduce the customer interface bandwidth, you can specify the `no-qos-adjust` statement at the `[edit routing-options multicast interface interface-name]` or `[edit dynamic-profiles profile-name routing-options multicast interface interface-name]` hierarchy level. When you configure the `no-qos-adjust` statement for the customer interface, available bandwidth is not reduced on the customer interface when multicast streams are added to the shared interface.



NOTE: Specifying the `no-qos-adjust` statement at the interface level also applies to the `reverse-oif-mapping` statement and is the preferred method for disabling QoS adjustment. The `reverse-oif-mapping no-qos-adjust` statement is deprecated but still supported.

4. (Optional) Define passive mode for IGMP or MLD.

A new `passive` statement exists at the `[edit protocols igmp interface interface-name]` and `[edit protocols mld interface interface-name]` hierarchy level. The `passive` statement specifies IGMP or MLD to run on the interface but to not send or receive control traffic (that is, IGMP or MLD reports, queries, and leaves).



NOTE: The OIF map interface should not typically pass IGMP or MLD control traffic and should be configured as passive. However, the OIF map implementation does support running IGMP or MLD on an interface (control and data) in addition to mapping data streams to the same interface. In this case, you should configure IGMP or MLD normally (that is, not in passive mode) on the mapped interface.

You can view OIF map-related information as follows:

- You can use the **show policy *policy-name*** command to view the OIF map as a route policy.
- The **show configuration protocols igmp interface *interface-name***, **show configuration protocols mld interface *interface-name***, **show igmp interface**, and **show mld interface** commands have been enhanced to display any OIF map association (when configured) and passive mode state (either *on* or *off*).
- The **show multicast interface** command has been enhanced to display the state of **no-qos-adjust** (when configured).
- New **show igmp output-group** and **show mld output-group** commands for IGMP and MLD display multicast group state for the interface that is sending the multicast data (that is, the output interface [OIF]). For interfaces with no OIF map, these new commands display very similar information as the **show igmp group** and **show mld group** commands because the control traffic interface is the same as the output interface. For interfaces with an OIF map, these new commands display the output interface specified by the OIF map. Adding the **detail** option adds an **Input interface** field to specify customer interfaces and a **Group mode** field to specify the group mode.
- **Translation of PIM join/prune messages to IGMP or MLD report/leave messages**—In some network configurations, customers are unable to run Protocol Independent Multicast (PIM) between the customer edge-facing PIM domain and the core-facing PIM domain, although PIM is running in sparse mode within each of these domains. Because PIM is not running between the domains, customers with this configuration cannot use PIM to forward multicast traffic across the domains. Instead, they might want to use Internet Group Management Protocol (IGMP) to forward IPv4 multicast traffic, or Multicast Listener Discovery (MLD) to forward IPv6 multicast traffic across the domains.

To enable the use of IGMP or MLD to forward multicast traffic across the PIM domains in such topologies, you can configure the rendezvous point (RP) router that resides between the edge domain and core domain to translate PIM join/prune messages received from PIM neighbors on downstream interfaces into corresponding IGMP or MLD report/leave messages. The router then transmits the report/leave messages by proxying them to one or two upstream interfaces that you configure on the RP router.

To configure the RP router to translate PIM join/prune messages into IGMP report/leave messages, include the **pim-to-igmp-proxy** statement at the **[edit routing-options multicast]** hierarchy level. Similarly, to configure the RP router to translate PIM join/prune messages into MLD report/leave messages, include the **pim-to-mld-proxy** statement at the **[edit routing-options multicast]** hierarchy level. As part of the configuration, you must specify the full name of at least one, but

not more than two, upstream interfaces on which to enable the PIM-to-IGMP proxy or PIM-to-MLD proxy feature.

To display the PIM-to-IGMP or PIM-to-MLD proxy state (enabled or disabled) and the name or names of the configured upstream interfaces, issue the **show multicast pim-to-igmp-proxy** command or the **show multicast pim-to-ml-proxy** command.

[*Multicast, Routing Protocols and Policies Command Reference*]

- **Flexible configuration for IGMP/MLD static-join**—When you configure static groups on an interface on which you want to receive multicast traffic, you can specify the number of static groups to be automatically created. Additionally you can specify an increment that controls the group addresses to be used when the static groups are automatically created.

To configure the number of static groups to be created, include the **group-count** statement and specify the number of groups to be created. To configure the group address increment, include the **group-increment** statement and specify the number by which the address should be incremented for each group. The increment is specified in dotted decimal notation similar to an IP address. For example, an increment of 0.0.0.3 increments the group address by three for each group created.

The **group-count** statement and **group-increment** statement can be configured at the [edit protocols igmp interface *interface-name* static group *multicast-group-address*] and [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address*] hierarchy levels.

When you configure a static group on an interface on which you want to receive multicast traffic, you can configure the number of source addresses associated with the static group to be automatically created. Additionally, you can specify an increment that controls the source addresses to be used when the static group source addresses are automatically created.

To configure the number of source addresses associated with a static group to be automatically created, include the **source-count** statement and specify the number of addresses to be created. To configure the source address increment include the **source-increment** statement and specify the number by which the address should be incremented for each source. The source increment is specified in dotted decimal notation similar to an IP address. For example, an increment of 0.0.0.2 increments the source address by two for each source created.

The **source-count** statement and **source-increment** statement can be configured at the [edit protocols igmp interface *interface-name* static group *multicast-group-address* source *ip-address*] and [edit logical-systems *logical-system-name* protocols igmp interface *interface-name* static group *multicast-group-address* source *ip-address*] hierarchy levels.

Similar configuration is available for IPv6 multicast traffic using the MLD protocol.

[*Multicast Protocols*]

- **Disabling PIM for IPv6 family**—Enables you to globally disable PIM for IPv6 exclusive of IPv4. Previously, you could only disable PIM for IPv4 and IPv6 simultaneously. This feature also lets you disable PIM on a particular IPv4 or IPv6 interface. Routing instances and logical interfaces are also supported.

To globally disable PIM for IPv4 or IPv6, include the **family *family-name* (disable)** statement at the [edit protocols pim] hierarchy level. To disable PIM for a particular

interface, include the `family family-name (disable)` statement at the `[edit protocols pim interface interface-name]` hierarchy level.

[Multicast Protocols]

Network Management

- **SNMP support for the TX Matrix Plus router**—To extend SNMP support to the newly-introduced TX Matrix Plus router, the integer values for `JnxChassisId` have been extended up to 28 from the previous maximum of 11. The updated textual convention for `JnxChassisId` is as follows:

```
unknown (1)
singleChassis (2),
scc (3),
lcc0 (4),
lcc1 (5),
lcc2 (6),
lcc3 (7),
jcs1 (8),
jcs2 (9),
jcs3 (10),
jcs4 (11),
node0 (12),
node1 (13),
sfc0 (14),
sfc1 (15),
sfc2 (16),
sfc3 (17),
sfc4 (18),
lcc4 (19),
lcc5 (20),
lcc6 (21),
lcc7 (22),
lcc8 (23),
lcc9 (24),
lcc10 (25),
lcc11 (26),
lcc12 (27),
lcc13 (28),
lcc14 (29),
lcc15 (30)
```

[Network Management]

- **SNMP support for Common Language Equipment Identifier (CLEI) code**—A new MIB object `jnxContentsChassisCleiCode` has been added to the Juniper Networks enterprise-specific Chassis MIB to store the CLEI code, also known as the hardware barcode, of the chassis.

A CLEI code is an intelligent code that consists of 10 alphanumeric characters with 4 data elements. The first data element is considered the basic code with the first two characters indicating the technology or equipment type, and the third and fourth characters denoting the functional subcategory. The second data element represents the features, and its three characters denote functional capabilities or changes. The third data element has one character and denotes

a reference to a manufacturer, system ID, specification, or drawing. The fourth data element consists of two characters and contains complementary data. These two characters provide a means of differentiating or providing uniqueness between the eight character CLEI codes by identifying the manufacturing vintage of the product. For more information about CLEI code, see http://www.commonlanguage.com/resources/commonlang/productshowroom/showroom/equip_id/carriers/overview.html.

[*Network Management*]

- **Enhancements to Juniper Networks enterprise-specific MPLS MIB**—The following objects of the enterprise-specific MPLS MIB (`jnx-mpls.mib`) have been modified to support and store information about manual bypass tunnels through the entire life cycle of a bypass tunnel.

Both `mplsLspState` and `mplsLspInfoState` objects now have two additional values: `notInService` (integer value: 4) and `backupActive` (5). The `notInService` state indicates that the LSP has been torn down or never been signaled due to the lack of demand for its protection. The `backupActive` state indicates that the LSP is up and carrying user traffic for at least one protected LSP due to the failure of the LSP, which has caused the creation of a backup LSP.

Similarly, the `mplsPathType` and `mplsPathInfoType` objects now have a new value, `bypass` (5), to denote that the path is a manually configured bypass tunnel.

In the previous releases, the information about bypass tunnels was stored in the standard `mplsTunnelTable` that uses a combination of `mplsTunnelIndex`, `mplsTunnelInstance`, `mplsTunnelIngressLSRId`, and `mplsTunnelEgressLSRId` as index. Because the value for `mplsTunnelInstance` changes when an LSP is signaled or resigaled, new entries are created each time an LSP is signaled or resigaled. This has been causing problems in tracking the state of bypass tunnels. The latest enhancements to the enterprise-specific MIB, which uses LSP name as index, enable the MIB to store information about bypass tunnels in a single entry and users to access information about bypass tunnels through its life cycle using a single index.

The `show mpls lsp bypass` command returns information about manually configured bypass tunnels of all states.

However, you are advised to:

- Set the `max-bypasses` value for RSVP interfaces to zero to disable creation of dynamic bypass tunnels.
- Assign unique names to the LSPs and bypass tunnels.

[*Network Management*]

Platform and Infrastructure

- **Routing Engine boot sequence for the TX Matrix Plus router and T1600 routers in a routing matrix**—The Routing Engines on the TX Matrix Plus router (or switch-fabric chassis) and T1600 routers (or line-card chassis) in the routing matrix boot from the storage media in this order: the USB device (if present), the CompactFlash card (if present), the disk (if present) in slot 1, and then the LAN.

[*TX Matrix Plus Hardware, Software Installation and Upgrade Guide, System Basics, System Basics and Services Command Reference*]

- **New TX Matrix Plus Router**—The TX Matrix Plus router is the centralized switch fabric of the routing matrix, which is a multi-terabit routing system for interconnecting T1600 routers. Each T1600 router connected to the TX Matrix Plus router adds 1.6 terabits per second (Tbps) of non-blocking subscriber switching capacity to the routing matrix. Currently, the TX Matrix Plus supports connections to up to four T1600 routers.



NOTE: TX Matrix Plus routers and T1600 routers that are configured as part of a routing matrix do not currently support nonstop active routing.

Routing Policy and Firewall Filters

- **Policer support for physical interfaces**—Enables you to define a policer for a physical interface and reference the policer in one or more firewall filters, which can then be applied to the physical interface. This feature enables you to configure a single aggregate policer for a physical interface that can be applied to all the logical interfaces and traffic families configured on the physical interface. Previously, you could define aggregate policers only for logical interfaces. A physical interface policer also enables you to apply a single policer to multiple routing instances because a physical interface policer includes all the logical interfaces configured on the physical interface even if they belong to different routing instances.

To configure a policer for a physical interface, include the **physical-interface-policer** statement at the [edit firewall policer *policer-name*] hierarchy level. To configure a firewall filter that references the physical interface policer, include the **physical-interface-filter** statement at the [edit firewall family *family-name* filter *filter-name*] hierarchy level. You must also apply the physical interface policer as an action for the firewall filter term. Include the **policer *policer-name*** statement at the [edit firewall family *family-name* filter *filter-name* term *term-name* then] hierarchy level. For *policer-name*, specify the name of a physical interface policer configured at the [edit firewall policer] hierarchy level. You can apply a firewall filter that references a physical interface policer as an input or output filter to any family on an interface. To apply a firewall filter as an input filter, include the **input *filter-name*** statement at the [edit interfaces *interface-name* unit *unit-number* family *family-name* filter] hierarchy level. To apply a firewall filter as an output filter, include the **output *filter-name*** statement at the [edit interfaces *interface-name* unit *unit-number* family *family-name* filter] hierarchy level.

The following example shows how to configure a physical interface policer and reference it in a firewall filter:

```
[edit firewall]
firewall {
  policer shared-police1 {
    physical-interface-policer;
    if-exceeding {
      bandwidth-limit 100m;
      burst-size-limit 500k;
    }
  }
}
```

```

    }
    then {
        discard;
    }
}
policer share-police2 {
    physical-interface-policer;
    if-exceeding {
        bandwidth-percent 5;
        burst-size-limit 250k;
    }
    then {
        discard;
    }
}
family inet {
    filter inet-filter {
        physical-interface-filter;
        term tcp-police-1 {
            from {
                precedence [ critical-ecp immediate-priority ];
                protocol tcp;
            }
            then policer shared-police1;
        }
        term tcp-police-2 {
            from {
                precedence [ internet-control routine ];
                protocol tcp;
            }
            then policer shared-police2;
        }
    }
}
}

```

[Policy]

- **Authentication for BFD (MD5/SHA1)**—Enables you to apply MD5 or SHA1 authentication requirements to BFD sessions on BGP, IPv4 static routes, IPv6 static routes, IS-IS, OSPFv2, PIM, and RIP protocols. Previously, BFD authentication was accomplished using a plain-text password. This feature also supports authentication key rollover for these protocols, which enables you to update authentication keys without causing associated neighboring sessions to reset. Authentication must be configured on both ends of the session. Additionally, you can configure loose authentication checking during a migration period from a nonauthenticated session to an authenticated session.

To enable MD5 or SHA1 authentication for BFD on an IPv4 or IPv6 static route, include the authentication algorithm *algorithm-name* statement and the authentication key-chain *keychain-name* statement at the [edit routing-options static route *address* bfd-liveness-detection] hierarchy level. You must also include the authentication-key-chains key-chain *keychain-name* statement at the [edit security] hierarchy level.

To enable MD5 or SHA1 authentication for BFD on BGP or RIP, include the authentication algorithm *algorithm-name* statement and the authentication key-chain *keychain-name* statement at the [edit protocols *protocol-name* group *group-name* bfd-liveness-detection] or [edit protocols *protocol-name* group *group-name* neighbor *address* bfd-liveness-detection] hierarchy level. You must also include the authentication-key-chains key-chain *keychain-name* statement at the [edit security] hierarchy level.

To enable MD5 or SHA1 authentication for BFD on IS-IS, OSPFv2, or PIM, include the authentication algorithm *algorithm-name* statement and the authentication key-chain *keychain-name* statement at the [edit protocols *protocol-name* interface *interface-name* bfd-liveness-detection] hierarchy level. You must also include the authentication-key-chains key-chain *keychain-name* statement at the [edit security] hierarchy level.

To enable graceful migration of nonauthenticated to authenticated sessions, include the authentication *loose-check* statement at the appropriate protocol hierarchy level.

You can configure this feature for all routing instances supported by BGP, IS-IS, OSPF, and RIP. Logical systems are also supported.

The `show bfd session detail` and `show bfd session extensive` commands have been enhanced to display authentication information when configured.

[*Routing Protocols, System Basics, Multicast Protocols, Routing Protocols and Policies Command Reference*]

- **Option to repartition jtree memory (M10i and M7i routers [with Enhanced CFEB], M320 router [with Enhanced III FPC1, Enhanced III FPC2, and Enhanced III FPC3], MX Series, and M120 Series routers)**—The jtree memory on I-CHIP-based Packet Forwarding Engines has two segments. One segment stores routing tables and related information. The other stores firewall-filter-related information. The new configuration statement `route-memory-enhanced` enables you to repartition the two jtree memory segments to allocate more memory for routing tables over firewall filters. This option is useful when you want to support larger routing tables. However, we recommend enabling this option only if you do not have a very large firewall filter configuration. To repartition more memory for routing tables, include the `route-memory-enhanced` statement at the [edit chassis] hierarchy level.

[*System Basics*]

- **Port mirroring for Layer 2 VPN (M120 and M320 routers)**—Enables port mirroring for Layer 2 VPNs on M120 and M320 routers, including firewall filtering capabilities. Previously, port mirroring was only supported for IPv4, IPv6, and VPLS traffic. This feature also enables you to set the maximum length of the mirrored packet.

To configure port mirroring for a Layer 2 VPN on an M120 or M320 router, include the `family ccc` statement at the [edit forwarding-options port-mirroring] hierarchy level. To configure port mirroring for a particular instance of a Layer 2 VPN, include the `family ccc` statement at the [edit forwarding-options port-mirroring instance *instance-name*] hierarchy level. Optionally, configure the maximum mirrored packet length using the `maximum-packet-length` statement at the [edit forwarding-options port-mirroring input] hierarchy level.

To configure a firewall filter for a Layer 2 VPN on an M120 or M320 router, include the `family ccc filter filter-name term term-name from match-conditions then action` statement at the `[edit firewall]` hierarchy level. You can apply the filter to a logical interface at the input or output. The filter can also be applied to aggregated Ethernet logical interfaces.

[Policy Framework, Network Interfaces]

Routing Protocols

- **Egress filtering PIMv4/v6 join messages**—You can filter PIM sparse mode (PIM-SM) join and prune messages at the egress interfaces for IPv4 and IPv6 in the upstream direction. The messages can be filtered based on the group address, source address, outgoing interface, PIM neighbor, or combination of these values.

This can be useful when the core of your network is using a mix of IP and MPLS. You can use this feature to selectively forward or filter PIM join and prune messages to PIM neighbors.

To filter PIM sparse mode join and prune messages at the egress interfaces, create a policy rejecting the group address, source address, outgoing interface, or PIM neighbor, then apply the policy.

To apply the policy, include the `export` statement at the `[edit protocols pim]`, `[edit routing-instances routing-instance-name protocols pim]`, `[edit logical-routers logical-router-name protocols pim]`, or `[edit logical-routers logical-router-name routing-instances routing-instance-name protocols pim]` hierarchy level.

You can display the number of messages that have been filtered using the `show pim statistics` command.

The PIM egress filter prevents PIM join and prune messages from being sent towards the upstream routers in PIM-SM and PIM source-specific multicast (PIM-SSM) protocol mode. This filter is not supported for PIM dense mode (PIM-DM).

[Multicast Protocols, Routing Protocols and Policies Command Reference]

- **Reduction in flooding of self-originated OSPF LSAs**—Enables you to override the default behavior of having OSPF flood self-originated link-state advertisements (LSAs) every 30 minutes. When this feature is enabled, the JUNOS Software floods self-originated LSAs with the `DoNotAge` bit set. As a result, LSAs are not reflooded unless a change occurs in the LSA. This feature reduces OSPF traffic overhead in stable topologies and facilitates the scaling of OSPF networks. You can enable flood reduction for OSPF interfaces, realms, virtual links, sham links, and peer interfaces. Additionally, you can configure this feature for all routing instances supported by OSPF. Logical systems are also supported.

To enable flood reduction for OSPF interfaces, include the `flood-reduction` statement at the `[edit protocols (ospf | ospf3) area area-id interface interface-name]` hierarchy level. To enable flood reduction for an OSPF realm, include the `flood-reduction` statement at the `[edit protocols ospf3 realm (ipv4-multicast | ipv4-multicast | ipv6-multicast) area area-id interface interface-name]` hierarchy level. To enable flood reduction for a virtual link, include the `flood-reduction` statement at the `[edit protocols (ospf | ospf3) area area-id virtual-link neighbor-id router-id transit transit-area]` hierarchy level. To enable flood reduction for a peer interface, include the `flood-reduction` statement at the `[edit protocols ospf area area-id`

`peer-interface interface-name]` hierarchy level. To enable flood reduction for the remote endpoint of a sham link, include the `flood-reduction` statement at the `[edit routing-instances routing-instance-name protocols ospf area area-id sham-link-remote address]` hierarchy level. The `show (ospf | ops3) interface extensive` command has been enhanced to display which interfaces are enabled for flood reduction.

[*Routing Protocols, Routing Protocols and Policies Command Reference*]

Services Applications

- **Border Gateway Function (BGF) RTCP sender and receiver reports (M120, M320, and T640 routers)**—Enables a BGF to send statistics for the RTCP sender and receiver reports to the gateway controller when the BGF receives an AUDIT or SUBTRACT request for a given gate.

To configure RTCP sender and receiver reports, include the `rtcp` statement at the `[edit services pgcp gateway gateway-name monitor media]` hierarchy level.

The `show services pgcp gateway gateway-name gate-id gate-id statistics` command now displays additional RTCP statistics: fraction lost, cumulative number of packets lost, and interarrival jitter.

[*Multiplay Solutions, Services Interfaces, System Basics and Services Command Reference*]

- **Integrated Multi-Service Gateway (IMSG) provides firewall and security services to SIP signaling traffic (M Series and MX Series routers)**—The IMSG provides firewall and security services to SIP signaling traffic before the traffic reaches the border signaling gateway (BSG). This feature uses a simple provisioning model where all the BSG protection elements are collected in a service set that is applied on a services interface subunit. After the service set processes the traffic, it is sent to the BSG. To configure, create a service set with the set of stateful firewall and IDP services that you wish to provide. You then apply the service set as an input service set and output service set on the interface at the `[edit interfaces interface-name unit logical-unit-number family inet service (input | output)]` hierarchy level.

[*Multiplay Solutions, Services Interfaces*]

- **Integrated Multi-Service Gateway (IMSG) support for manipulating SIP headers (M120, M320, and T640 routers)**—You can now add, remove, or change the value of headers in incoming SIP messages. This feature provides greater interoperability by modifying SIP messages so that they can pass between different peers and vendors. To configure, create a message manipulation rule at the `[edit services border-signaling-gateway gateway gateway-name sip message-manipulation-rules]` hierarchy level. You then create actions in new transaction policies that specify the message manipulation rules that are applied to either incoming or outgoing messages. Configure these actions at the `[edit services border-signaling-gateway gateway gateway-name sip new-transaction-policy policy-name term term-name then message-manipulation]` hierarchy level. To display the message manipulation rules that are in effect for a contact or request URI, use the detail level of the `show services border-signaling-gateway by-contact` or `show services border-signaling-gateway by-request-uri` operational mode commands.

[*Multiplay Solutions, Services Interfaces, System Basics and Services Command Reference*]

- **ping mpls command enhancements**—Enable you to specify the packet size for LSP ping messages transmitted when MPLS-based services are deployed. LSP ping enables providers of these services to diagnose network problems. Previously, the packet size could not be user-defined. This feature enables configuration of the LSP ping packet size for Layer 2 circuits, Layer 2 VPNs, Layer 3 VPNs, LDP, LSP endpoints, and RSVP services. Using this feature you can determine the size of the MTU for LSP ping messages either manually or automatically. This information aids in troubleshooting misconfigured MTUs for an LSP that might cause errors on the network.

To specify the LSP ping packet size for LDP, RSVP, Layer 3 VPN, and LSP, execute the `ping mpls service-name size` command. To enable automated MTU size determination for these services, execute the `ping mpls service-name sweep` command. To specify the LSP ping packet size for Layer 2 circuits at the interface or virtual circuit level, execute the `ping mpls l2circuit (interface | virtual-circuit) size` command. To enable automated MTU size determination for this service, execute the `ping mpls l2circuit (interface | virtual-circuit) sweep` command. To specify the LSP ping packet size for Layer 2 VPNs at the interface or instance level, execute the `ping mpls l2circuit (interface | instance) size` command. To enable automated MTU size determination for this service, execute the `ping mpls l2circuit (interface | instance) sweep` command.

[*System Basics and Services Command Reference*]

- **Dynamic application awareness functionality extended to M320 routers**—Adds support for Intrusion Detection and Prevention (IDP) and application identification (APPID) functionality on M320 routers equipped with MultiServices (MS100 and MS400) PICs. The feature set is the same as documented for MX Series MultiServices DPCs; there are no new CLI statements or commands. The fast update filter (FUF) feature is not supported on M320 routers in this release.

[*J Series Services Router Guides, Services Interfaces*]

- **Border gateway function (BGF) usability and H.248 compliance (M120, M320, and T640 routers)**—This release provides the following usability and H.248 compliance improvements:
 - Publish supported packages—A virtual BGF can now publish a list of its supported packages, including their version.
 - Always return valid SDP in replies—SDPs returned by the virtual BGF are fully valid. The BGF does not allow omission of any mandatory fields.
 - Enforce the negotiated H.248 version—The virtual BGF now confirms for messages it sends that the version of the message header matches the negotiated version and that the content of the message is valid in the negotiated version. However, the BGF does accept messages whose content does not match the version of the header.
 - Enforce range (minimum and maximum values) for CLI parameters—In addition to the default value configured to H.248 parameters, the operator can configure the minimum and maximum values supported for specific parameters. The range of parameter values has predefined minimum and maximum values according to the BGF functionality, but the operator can use the CLI to narrow the range. The corresponding parameters received via the H.248 interface should be in the range configured in CLI. If not, an H.248

error code #449: "Unsupported or Unknown Parameter or Property Value" should be issued.

- Prevent TCP SendOnly and RecvOnly—The stream modes SendOnly and RecvOnly are meaningless for TCP (TCP ACKs must always flow). When an attempt is made to create a TCP stream with SendOnly or RecvOnly, the BGF now returns error #473: "Conflicting property values."
- Support encoding of H.248 messages in lowercase—By default H.248 messages are encoded in uppercase which can be changed to lowercase using the following command: `edit service pgcp gateway gw h248-options encoding use-lower-case`.
- **Exclude media data inactivity for RTCP streams**—By default, media inactivity is detected for both RTP and RTCP streams. An RTCP stream can be excluded using the following command: `edit service pgcp gateway gw data-inactivity-detection no-rtcp-check`.

[*Multiplay Solutions, Services Interfaces, System Basics and Services Command Reference*]

- **Port mirroring with next-hop groups enhancement**—Adds support for port mirroring functionality using next-hop groups without Tunnel PICs on MX Series routers. Previously, port mirroring using next-hop groups, also known as multipacket port mirroring, was supported on MX Series routers, but required installation of a Tunnel PIC.

To configure the new functionality, include the `next-hop-group` statement at the `[edit forwarding-options port-mirror family inet output]` or `[edit forwarding-options port-mirror instance instance-name family inet output]` hierarchy level. You can disable this configuration by including a `disable` or `disable-all-instances` statement at the `[edit forwarding-options port-mirror]` hierarchy level or by including a `disable` statement at the `[edit forwarding-options port-mirror instance instance-name]` hierarchy level. You can display the settings and network status by issuing the `show forwarding-options next-hop-group` and `show forwarding-options port-mirroring` operational commands. [*Services Interfaces, System Basics and Services Command Reference*]

- **Active flow monitoring sampling enhancements**—Add support for packet sampling on a per Packet Forwarding Engine basis. You can now configure active sampling by defining a sampling instance that specifies a name for the sampling parameters and binding the instance name to a particular Packet Forwarding Engine. To specify the sampling parameters, include the `instance` statement at the `[edit forwarding-options sampling]` hierarchy level. To bind an instance to a Packet Forwarding Engine, include the `sampling-instances` statement at the `[edit chassis fpc number]` hierarchy level. You can also set global input parameters by including the `rate` and `run-length` statements at the `[edit forwarding-options sampling input]` hierarchy level. In a related development, a proprietary v5 extension template is now available for supporting 4-byte AS information in flow records. To configure this extension, include the `version 5` statement at the `[edit forwarding-options sampling output flow-server server-name]` hierarchy level. [*Services Interfaces*]
- **Support for MultiServices PICs (JCS1200 platform)**—MultiServices 400 PICs and IPSec are supported in this release.

[PSD Configuration Guide]

- **Border Gateway Function (BGF) support for pgcpd process running on MS-PIC or MS-DPC (M120, M320, and T640 routers)**—The pgcpd process for the BGF can now run on either the Routing Engine or on an MS-PIC or MS-DPC. You can configure up to eight virtual BGFs on the PIC or DPC. You must enable the BGF service package on the PIC or DPC by including the `set fpc slot pic slot adaptive-services service-package extension-provider package jservices-bgf-pic` statement at the `[edit chassis]` hierarchy level. To specify whether the virtual BGF runs on the PIC or DPC or on the Routing Engine, include the `platform` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level. All `show services pgcp` operational mode commands now allow you to display either a backup or master PIC or DPC for a virtual BGF.

[Multiplay Solutions, System Basics and Services Command Reference]

- **Border Gateway Function (BGF) support for high availability for PICs running the pgcpd process (M Series and T Series routers with MS-PICs and MX Series routers with MS-DPCs)**—The BGF provides a 1:1 redundancy model for PICs and DPCs that are running the pgcpd process. There is one primary service PIC and one secondary service PIC that acts as a backup. If the primary service PIC fails, the secondary PIC becomes active and H.248 traffic is switched over to the new active PIC. When the failed PIC recovers, it comes up as the secondary PIC. The BGF supports hot standby mode to ensure that failover occurs in 5 seconds or less.

To configure, create a container interface called an rms interface at the `[edit interfaces]` hierarchy level. On the rms interface, configure the primary and secondary PICs or DPCs. When you configure your virtual BGF, specify the rms interface as the platform device on which the virtual BGF runs by including the `platform` statement at the `[edit services pgcp gateway gateway-name]` hierarchy level.

All `show services pgcp` operational mode commands now allow you to display information about either the backup or the master PIC or DPC for a virtual BGF.

[Multiplay Solutions, Services Interfaces, System Basics and Services Command Reference]

- **Border Gateway Function (BGF) RTCP bandwidth policing (M120, M320, and T640 routers)**—Enables a gateway controller to control RTCP bandwidth using two RTCP bandwidth modifiers in the H.248 stream. The “RR” modifier specifies the SDR (sustained data rate) bandwidth for active receivers, while the “RS” modifier specifies the SDR bandwidth for active senders. When the H.248 stream does not contain the RTCP bandwidth modifiers, the bandwidth specified by the CLI configuration is used.

You can now specify whether to include RTCP bandwidth in the traffic management SDR for *all streams*. To include RTCP bandwidth, include the `rtcp-include` statement at the `[edit services pgcp gateway gw-name h248-properties traffic-management sustained-data-rate]` hierarchy level.

[Multiplay Solutions, Services Interfaces]

- **Fast update filters for dynamic profiles (MX Series routers)**—Fast update filters are now supported for dynamic profiles.

Fast update filters are a unique type of firewall filter that enable you to incrementally add, remove, or update the filter terms. Dynamically assigned fast update filters provide significant advantages over static filters. Networks can have multiple subscribers who might be mapped to a single logical interface. Because static filters are compiled at commit time, they cannot contain subscriber-specific terms, and therefore cannot distinguish between different subscribers.

Using fast update filters involves two procedures. The two procedures are the same as the existing procedures used for normal filters. You first configure the filter, and you then use a dynamic profile to apply the filter to an interface family.

- To configure fast update filters, include the **fast-update-filter** statement at the [edit dynamic-profiles profile-name firewall family inet] hierarchy level. When used with dynamic profiles, fast update filters are interface-specific and the order of the match field is explicit. You must include the **interface-specific** statement and the **match-order** statement when you configure a fast update filter.
- To apply a fast update filter to an interface family, include the **filter** statement at the [edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family inet] hierarchy level.

When a dynamic profile instantiates a session, the router verifies if the applied fast update filter is present on the session's interface. The router matches filters by name.

- If a filter with the same name is not present, the router adds all terms of the new filter.
- If a filter with the same name is present, the router adds only new or modified terms to the filter.



NOTE: All terms that you add to an existing filter must be unique. You can ensure the required uniqueness by using the `$junos-subscriber-ip-address` variable as either the **source-address** (input filter) or **destination-address** (output filter) in the **from** statement at the [edit dynamic-profiles profile-name firewall family inet fast-update-filter term] hierarchy level.

Use the **show firewall** and **show firewall filter** operational commands to view information for fast update filters.

[Subscriber Access, Policy Framework]

- **New option for the show snmp mib command**—The **show snmp mib** command now has a new option **ascii** that converts the string indices to an “ascii-key” representation, thus making the output more readable. This is in addition to the default option of the decimal format. The syntax for the command is **show snmp mib (get | get-next | walk) (ascii | decimal) object-id**.

[System Basics and Services Command Reference]

- **Integrated Multi-Service Gateway (IMSG) support for the following features (M120, M320, MX Series, and T640 routers):**
 - **NAT pool selection**—You can now assign different NAT pools for different networks (or peers) and for different network subnets. You do so by assigning a virtual interface to the ingress and egress service points. The virtual interface specifies the NAT pool to be used on the service point. To configure, include the `default-media-realm` statement at the `[edit services border-signaling-gateway gateway gateway-name service-point service-point-name]` hierarchy level. The value of the `default-media-realm` is a virtual interface that you configured at the `[edit services pgcp]` hierarchy level.
 - **Media inactivity detection**—The border signaling gateway (BSG) now supports media inactivity detection. To configure, include the `data-inactivity-detection` statement at the `[edit services border-signaling-gateway gateway gateway-name sip new-call-usage-policy policy-name term term-name then media-policy]` hierarchy level. To discontinue media inactivity detection while a call is on hold, include the `stop-detection-on-drop` statement at the `[edit services pgcp gateway gateway-name data-inactivity-detection]` hierarchy level.
 - **Media anchoring**—You can now explicitly turn media anchoring on or off. Media anchoring is on by default. To turn media anchoring on or off, use the `no-anchoring` statement at the `[edit services border-signaling-gateway gateway gateway-name sip new-call-usage-policy v term term-name then media-policy]` hierarchy level.
 - **DSCP**—The DSCP value in the embedded SPDF is changed to a 6-bit value instead of an 8-bit value. If you do not specify a DSCP value, the default value is `do-not-change`. Setting the DSCP to a hexadecimal value is no longer supported. To configure DSCP values, include the `dscp` statement at the `[edit services border-signaling-gateway gateway gateway-name embedded-spdf service-class service-class-name term term-name then]` hierarchy level.

[Multiplay Solutions, Services Interfaces]

- **New statement for controlling session offload behavior (MX Series routers)**—The new `session-offload` statement at the `[edit chassis fpc slot-number pic number adaptive-services service-package extension-provider]` hierarchy level controls session offload behavior for MS-DPCs on MX Series routers. It controls session offload on a per-device basis, where a device is a MultiServices interface (`ms-fpc-pic-port`). In JUNOS Release 9.6, the session offload function is supported for at most one MultiServices interface. When the offload function is enabled, it is strongly recommended that you limit Dynamic Application Awareness features to that MultiServices interface. The default is to not offload any sessions.

[System Basics]

Subscriber Access Management

- **JUNOS subscriber access scaling values**—The following subscriber access scaling values are supported in this release:
 - Number of subscriber VLANs per DPC: 16,000
 - Number of subscriber VLANs per chassis for the MX240 router, which accommodates 2 DPCs: 32,000
 - Number of subscriber VLANs per chassis for the MX480 and the MX960 routers: 64,000
 - Number of DHCP bindings: 120,000
- **Diameter base protocol support for Authentication, Authorization, and Accounting (AAA) (MX Series routers)**—The Diameter protocol is defined in *RFC 3588, Diameter Base Protocol*, and provides an alternative to RADIUS that is more flexible and extensible. The Diameter base protocol provides basic services to one or more applications (also called functions) that each runs in a different Diameter instance. The individual application provides the extended AAA functionality. JSRC is the first application that is supported.

Diameter peers communicate over a TCP transport layer connection by exchanging Diameter messages that convey status, requests, and acknowledgments by means of standard Diameter AVPs and application-specific AVPs. The Diameter transport layer configuration is based on Diameter network elements (DNEs); multiple DNEs per Diameter instance are supported. Currently only the predefined *master* Diameter instance is supported, but you can configure alternative values for many of the master Diameter instance values.

Each DNE consists of a prioritized list of peers and a set of routes that define how traffic is forwarded. Each route associates a destination with a function, a function partition, and a metric. When an application sends a message to a routed destination, all routes within the Diameter instance are examined for a match. When the best route to the destination has been selected, the message is forwarded by means of the DNE that includes that route.

To configure a Diameter network element, include the **network-element** statement at the [edit diameter] hierarchy level. Include the **route** statement at the [edit diameter network-element *element-name* forwarding] hierarchy level. To configure a route for the DNE, include the **destination** (optional), **function** (optional), and **metric** statements at the [edit diameter network-element *element-name* forwarding route *dne-route-name*] hierarchy level. Specify the Diameter peers associated with the DNE by including one or more **peer** statements at the [edit diameter network-element *element-name*] hierarchy level. Set the priority for each peer with the **priority** statement at the [edit diameter network-element *element-name* peer *peer-name*] hierarchy level.

Diameter requires you to configure information about the origin node; this is the endpoint node that originates Diameter for the Diameter instance. Include the **host** and **realm** statements at the [edit diameter] hierarchy level to configure the Diameter origin.

Each Diameter peer is specified by a name. Peer attributes include address and the destination TCP port used by active connections to this peer. To configure a

Diameter peer, include the **peer** statement at the [edit diameter] hierarchy level. Include the **address** and **connect-actively** statements at the [edit diameter peer *peer-name*] hierarchy level. To configure the active connection, include the **port** statement at the [edit diameter peer *peer-name* connect-actively] hierarchy level.

[Subscriber Access]

- **Support for dynamic 802.1Q VLAN configuration over aggregated Ethernet interfaces**—You can now configure dynamic 802.1Q VLANs over aggregated Ethernet interfaces.

[Subscriber Access]

- **DHCPv6 local server (MX Series routers)**—Enables the router to function as a server for the DHCP protocol for IP version 6 (IPv6). The DHCPv6 local server sends and receives packets using the IPv6 protocol and informs IPv6 of the routing requirements of the router clients. The local server works together with the AAA service framework to control subscriber access and accounting. For this release, DHCPv6 local server uses RADIUS to supply the IPv6 prefix and client configuration parameters. As a prerequisite for using the DHCPv6 local server with RADIUS, you must configure DHCPv6 authentication.

The DHCPv6 local server supports the following RADIUS attributes and vendor-specific attributes (VSAs):

- Session-Timeout (attribute 27)
- NAS-IPv6-Address (attribute 95)
- Framed-Interface-Id (attribute 96)
- Framed-IPv6-Prefix (attribute 97)
- Login-IPv6-Host (attribute 98)
- Framed-IPv6-Pool (attribute 100)
- Delegated-IPv6-Prefix (attribute 123)
- Maximum-Clients-Per-Interface (VSA 26-143)

To configure a DHCPv6 local server, use the **dhcpv6** statement at the [edit system services dhcp-local-server] hierarchy level. You can configure DHCPv6 globally or for a named group of interfaces.

The **dhcpv6 authentication username-include** statement, which specifies the **username** that is passed to AAA, supports the following new parameters, in addition to the parameters supported by the IPv4 DHCP local server:

- **relay-agent-remote-id**—The DHCPv6 Relay Agent Remote-ID option (option 37) from the client PDU name.
- **relay-agent-subscriber-id**—The DHCPv6 Relay Agent Subscriber-ID option (option 38) from the client PDU name.

To view DHCPv6 configuration and information, use the **show dhcp server bindings** command.

DHCPv6 local server is compatible with the extended DHCP local server and DHCP relay. DHCPv6 local server does not support dynamic profiles in this release.

[*Subscriber Access*]

- **Support for DHCP Layer 3 wholesale configuration in a subscriber access network**—Enables you to configure DHCP Layer 3 wholesaling within a subscriber access network. Wholesale access is the process by which an access network provider partitions the access network into separately manageable and accountable subscriber segments for resale to other network providers. An access network provider might elect to wholesale all or part of its network to one or more service providers (retailers).

In a Juniper Networks subscriber access network, you accomplish Layer 3 partitioning through the use of logical systems and routing instances. Logical systems enable you to divide the physical router into separate, distinct, logical administrative domains. This method of division enables multiple providers to administer the router simultaneously and each have access to only the portions of the configuration that are relevant to their specific logical system. JUNOS software supports up to 15 named logical systems in addition to the default logical system (inet.0).

Routing instances are typically used in Layer 3 VPN scenarios. A routing instance does not have the same level of administrative separation as does a logical system. The routing instance defines a distinct routing table, set of routing policies, and set of interfaces, but it does not provide administrative isolation.



NOTE: In this release, DHCP Layer 3 wholesaling supports the use of only the default logical system using multiple routing instances.

When configuring DHCP Layer 3 wholesale for a subscriber access network, keep the following in mind:

- Routing instances must use the same DHCP configuration: DHCP Local Server or DHCP Relay.
- The configuration must include at least a single address pool with a single range of addresses. However, the configuration can include multiple address pools as long as each address pool uses non-overlapping address ranges.
- You cannot configure the **use-primary** statement within a DHCP Layer 3 wholesale configuration. Configuring the **use-primary** statement on an underlying logical interface results in no routing instance reassignment when a client requiring redirection to a retail routing instance logs in. The client is subsequently logged out.
- Each routing instance must contain a loopback with one or more addresses to be used for the unnumbered interface. The loopback must also have an address within the subnet of each possible client IP address.

To configure DHCP Layer 3 wholesale for a subscriber access network:

- Include the `routing-instances` statement along with the `$junos-routing-instance` dynamic variable at the `[edit dynamic-profiles profile-name interface $junos-interface-name]` hierarchy level.
- Include the `interface` statement along with the `$junos-interface-name` dynamic variable at the `[edit dynamic-profiles profile-name interface "$junos-interface-name" routing-instances "$junos-routing-instance"]` hierarchy level.
- Include the `unnumbered-address` statement along with the `$junos-loopback-interface` dynamic variable along with the `preferred-source-address` option and the `$junos-preferred-source-address` dynamic variable at the `[edit dynamic-profiles profile-name interfaces demux0 unit "$junos-interface-unit" family inet]` hierarchy level.

To view the logical system and routing instance for each subscriber, use the `show subscriber operational` command.

- **Enhanced PPP support (M120 and M320 routers)**—IP address hinting for PPPoE is now supported. Consequently, you can use the PPPoE destination address as the assigned peer address. You specify this address by including the `destination` statement at the `[edit interfaces interface-name family inet address address]` hierarchy level.

The no-authentication configuration for PPPoE is also now supported.

These two feature enhancements remove the requirement that you must always have a RADIUS server for subscriber-aware services for PPP sessions.

[Subscriber Access]

- **JSRC supports subscriber sessions on static and dynamic interfaces (MX Series routers)**—You can use the new JSRC application to interact with the Juniper Networks Session and Resource Control (SRC) software to support dynamic and static subscribers. The SRC software runs on a Juniper Networks C Series Controller and provides a central administrative point for managing subscribers and their services. The SRC software uses the Diameter protocol for communications between JSRC acting as the local SRC peer on an MX Series router and the remote SRC peer (the service activation engine or SAE) on a C Series Controller.

JSRC requests address and service authorizations from the SAE, activates and deactivates services as specified by the SAE, logs out subscribers as specified by the SAE, and synchronizes subscriber state and service information with the SAE.

JSRC is a Juniper Networks-specific Diameter application registered with the IANA as Juniper-Policy-Control-JSRC, with an ID of 16777244. JSRC and the SAE exchange Diameter protocol messages that include a variety of attribute-value pairs (AVPs) to convey state information and identify actions requested or performed. Both standard Diameter AVPs and Juniper Networks vendor-specific AVPs (ID 2636) are employed.

JSRC configuration includes creating and assigning a JSRC partition, enabling authorization of DHCP subscribers, and enabling SRC provisioning of subscriber services.

To create the JSRC partition, include the **partition** statement at the [edit jsrc] hierarchy level. To configure the partition, include the **diameter-instance**, **destination-host**, and **destination-realm** statements at the [edit jsrc partition *partition-name*] hierarchy level. To assign the JSRC partition, include the **jsrc-partition** statement at the [edit] hierarchy level.

To enable SRC authorization of DHCP subscribers, include the **authorization-order jsrc** statement at the [edit access profile *profile-name*] hierarchy level.

To enable SRC provisioning for DHCP subscribers, include the **provisioning-order jsrc** statement at the [edit access profile *profile-name*] hierarchy level.

[Subscriber Access]

- **Dynamic subscriber and service management on static interfaces (MX Series routers)**—You can now associate subscribers with statically configured interfaces and provide dynamic service activation for these subscribers. When the interface comes up, it is treated as a subscriber login. When the interface goes down, it is treated as a subscriber logout.

Static subscribers are intended to work with JSRC. After the subscribers are present in the session database (SDB), JSRC can report the subscribers to the SAE so that SRC software can subsequently manage the subscribers. Alternatively, you can configure the static subscribers to be authenticated and authorized by means of RADIUS. In this case, RADIUS can then activate and deactivate services with change of authorization (CoA) messages. However, this configuration does not prevent the interface from coming up and forwarding traffic. Further, authorization parameters are not imposed on the subscriber interface.

Currently, only Ethernet interfaces support static subscribers. Only one static subscriber can exist over a given interface. An interface cannot appear in more than one group. Static subscribers cannot be created over dynamic interfaces.

To enable JSRC to handle the subscribers at the direction of the SRC software, include the **provisioning-order jsrc** statement at the [edit access profile *profile-name*] hierarchy level. If the authentication request fails for a static subscriber, the request is reissued after one hour, when a nonconfigurable timer expires. This action repeats for as long as the interface is operationally up.

You can force a logout of the static subscriber by issuing the **request services static-subscribers logout interface *interface-name*** command. A static subscriber can also be logged out by AAA or an external policy manager. In both cases, no subsequent logins can take place on the underlying interface until you reset the state by issuing the **request services static-subscribers login interface *interface-name*** command or the router or process reboots.

When you commit configuration changes for existing interface groups, static subscribers are logged out and then back in. You can force a logout of an interface group by issuing the **request services static-subscriber logout group *group-name*** command. You can subsequently log in a group of interfaces by issuing the **request services static-subscriber login group *group-name*** command.

No new CLI statements are required to configure the dynamic profile for static subscribers. The client profile can be very simple; it is activated at login and deactivated at logout. During a graceful Routing Engine switchover (GRES) event, active static subscribers are recovered, inactive subscribers are cleaned up, and logout continues for subscribers that were in the process of logging out.

To configure static subscribers, include the **static-subscribers** statement at the **[edit system services]** hierarchy level. To configure tracing operations for static subscribers, include the **traceoptions** statement at the **[edit system services static-subscribers]** hierarchy level. You can configure the access profile, dynamic profile, and authentication parameters for all groups or for a particular group.

To configure the access profile that triggers AAA services for the static subscriber for all subscribers, include the **access-profile** statement at the **[edit system services static-subscribers]** hierarchy level. Alternatively, include the statement at the **[edit system services static-subscribers group group-name]** hierarchy level to apply the profile to a specific group and override a top-level configuration.

To configure the dynamic profile that is instantiated when the static subscriber logs in for all subscribers, include the **dynamic-profile** statement at the **[edit system services static-subscribers]** hierarchy level. Alternatively, include the statement at the **[edit system services static-subscribers group group-name]** hierarchy level to apply the profile to a specific group and override a top-level configuration. Do not specify a dynamic profile that creates a dynamic interface.

To configure the authentication parameters that trigger an Access-Request message to AAA for all subscribers, include the **authentication** statement at the **[edit system services static-subscribers]** hierarchy level. Alternatively, include the statement at the **[edit system services static-subscribers group group-name]** hierarchy level to configure authentication for a specific group and override a top-level configuration. If you do not configure authentication, then by default the interface name is modified and used as the default username for the client session and the authentication request.

The configurable authentication parameters include the password and details of how the username is formed. To configure the authentication password for all subscribers, include the **password** statement at the **[edit system services static-subscribers authentication]** hierarchy level. Alternatively, include the statement at the **[edit system services static-subscribers group group-name authentication]** hierarchy level to configure authentication for a specific group and override a top-level configuration.

The username that is sent to AAA for authentication must include at least one of the following attributes: the domain name, a user prefix, the interface name, a logical system name, or a routing instance name. To configure how the username is formed for all subscribers, include the desired statements at the **[edit system services static-subscribers authentication]** hierarchy level: **domain-name**, **user-prefix**, **logical-system-name**, or **routing-instance-name**. Alternatively, include the desired statements at the **[edit system services static-subscribers group group-name authentication]** hierarchy level to configure the username for a specific group and override a top-level configuration.

Finally, a group configuration must specify all the interfaces that you expect to support static subscribers. You must also statically configure these interfaces before any static subscribers can be supported on them. To specify the interfaces, include the **interface** statement at the **[edit system services static-subscribers group group-name]** hierarchy level. This statement enables you to specify a single interface or a range of interfaces.

To display information about the state of static subscribers, you can issue any of the following commands:

- `show static-subscribers`
- `show static-subscribers interface interface-name`
- `show static-subscribers group group-name`

The following new system log messages are supported:

- `JSSCD_AUTH_CONNECTION_FAILURE`
- `JSSCD_DAX_REGISTER_FAILURE`
- `JSSCD_DYN_PROFILE_FAILURE`
- `JSSCD_INIT`
- `JSSCD_RTsock_OPEN_FAILURE`
- `JSSCD_RTsock_REGISTER_FAILURE`
- `JSSCD_SDB_OPEN_FAILED`
- `JSSCD_TRACE`

[Subscriber Access]

- **Using the textual interface description in DHCP relay option 82 Agent Circuit ID suboption (MX Series routers)**—DHCP relay agent option 82 support enables you to include the textual interface description in the Agent Circuit ID suboption for static interfaces. By default, when DHCP option 82 is inserted into client packets, the Agent Circuit ID suboption includes the interface identifier. You can now optionally specify that the Agent Circuit ID suboption include the textual description that is configured for the interface instead of the interface identifier. Using the textual description can provide a more descriptive and flexible identification as compared to the standard circuit identifier that contains router interface information.

To specify that the Agent Circuit ID suboption use the textual interface description, configure the `use-interface-description` statement at the `[edit forwarding-options dhcp-relay relay-option-82 circuit-id]` hierarchy level. You can specify that the Agent Circuit ID use the textual description for either the logical interface or the device interface.

[Subscriber Access]

User Interface and Configuration

- **CLI changes for the TX Matrix Plus router**—JUNOS Software support for the TX Matrix Plus router introduces some new CLI options. The `sfc` option introduced in the operational mode commands and configuration statements represents the TX Matrix Plus router (or switch-fabric chassis). The `lcc` option is now also supported for T1600 routers (or line-card chassis) in a routing matrix.

In addition, the **all-chassis** and **all-lcc** options are now also supported for some commands on the TX Matrix Plus router.

The **all-sfc** option is introduced in JUNOS Release 9.6. However, it is reserved for future configurations that can support more than one TX Matrix Plus router (or switch-fabric chassis) in a routing matrix. Currently, only one TX Matrix Plus router is supported in a routing matrix.

The **lcc number** and **sfc** options have been added to the **request system software add**, **request routing-engine login**, and **file copy** and **file rename** commands to support installation of a software package or bundle on a TX Matrix Plus router or any of the T1600 routers in the routing matrix. The **sfc** option is used to install the software package or bundle on a TX Matrix router. The **lcc number** option is used to install a software package or bundle on a T1600 router in the routing matrix.

[TX Matrix Plus Hardware Guide, Software Installation and Upgrade Guide, System Basics and Services Command Reference]

- **Remote RPC support for event scripts**—Enables you to remotely execute RPCs from event scripts. To configure support for RPC execution, include the remote hostname(s) at the new hierarchy level **[edit event-options event-script file filename remote-execution]**. You must configure the **username** and **passphrase** statements for each of the remote hosts at the **[edit event-options event-script file filename remote-execution remote-hostname]** hierarchy level. The remote hostname, username, and passphrase, in addition to the event details, are passed to the event script when it is triggered by an event policy. You can configure the event script to establish a connection with the remote host by invoking the enhanced **jcs:open()** function, which now accepts a username and passphrase for a remote device. You can then execute the RPCs from within the event script with the **jcs:execute()** extension function.

[Configuration and Diagnostic Automation Guide]

- **New login script feature**—Enables you to configure the router to automatically execute an op script when a user in a designated user class logs in to the CLI. To associate an op script with a given class, include the **login-script** statement and specify the filename of the op script at the **[edit system login class class-name]** hierarchy level. The configured op script is executed when any user that has the configured **class class-name** logs in to the CLI, provided that the script has been enabled.

[Configuration and Diagnostic Automation Guide]

VPNs

- **GRE tunnels for Layer 3 VPNs**—You can now configure a GRE tunnel from a PE router to a remote CE router in a Layer 3 VPN. Previously, you could only configure a GRE tunnel from a PE router to a P router or from a PE router to a local CE router. Configure the `source` statement, the `destination` statement, and the `routing-instance` statement at the `[edit interfaces gr-fpc|pic|port unit unit-number tunnel]` hierarchy level to configure the GRE tunnel on the PE router. Configure the `source` statement and the `destination` statement at the `[edit interfaces gr-fpc|pic|port unit unit-number tunnel]` hierarchy level to configure the GRE tunnel on the remote CE router.

[VPNs]

- **Connectivity-fault management process flooding to interfaces based on mesh groups**—A mesh group is a set of interfaces with similar flooding behavior, commonly used in the configuration of LDP BGP VPLS interworking. The flooding behavior configured for each mesh group is used by all the interfaces in that mesh group. The Connectivity Fault Management Protocol uses the data path semantics for packet flooding and therefore must adhere to the flooding behavior for mesh groups. The connectivity-fault management process (cfmd) can now support VPLS and virtual-switch routing instances.

[VPNs]

- **Load balancing and IP header filtering for Layer 3 VPNs**—You can now simultaneously enable both load balancing of traffic across both internal and external BGP paths and filtering of traffic based on the IP header. To enable these features, include the `vpn-unequal-cost equal-external-internal` statement at the `[edit routing-instances routing-instance-name routing-options multipath]` hierarchy level and the `vrf-table-label` statement at the `[edit routing-instances routing-instance-name]` hierarchy level.

[VPNs]

- **PIM source-specific multicast (PIM-SSM) provider tunnel support added to Multiprotocol BGP-based multicast VPNs: next-generation**—This feature allows service providers to configure PIM-SSM in the core network to carry customer data. PIM-SSM tunnels can be configured as an inclusive provider multicast service interface (PMSI).

In addition, a new configuration statement is being added to control the single forwarder election, as described in the Internet draft *draft-ietf-l3vpn-2547bis-mcast-bgp*.

By default, the single forwarder election is determined by the IP address from the global administrator field in the `route-import` community.

You can configure a router to use the unicast route preference to determine the single forwarder election. To configure the router to use the unicast route preference, include the `unicast-umh-election` statement. The `unicast-umh-election` statement can be configured at the `[edit logical-systems logical-system-name routing-instances routing-instance-name Protocols mvpn]` and `[edit routing-instances routing-instance-name protocols mvpn]` hierarchy levels.

[VPNs, Multicast]

- Related Topics**
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 46
 - Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 54
 - Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers on page 96
 - Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 103

Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers

Forwarding and Sampling

- **Enhancement to the show firewall command**—The show firewall command now supports a terse option that enables you to display only the names of firewall filters. This option displays no other information about the firewall filters configured on your system. Use the show firewall terse command to verify that all the correct filters are installed.

[Routing Protocols and Policies Command Reference]

Interfaces and Chassis

- **Deactivation and inheritance differences between static and dynamic interfaces**—When you deactivate a statically created physical interface, the logical interfaces (if any) contained within that interface is deactivated. However, for a dynamic interface, if you deactivate the physical interface, then the logical interfaces contained within that interface continues to forward the traffic. In addition, the dynamically created logical interfaces do not inherit the changes made to that physical interface.

When using dynamically created logical interfaces, delete or modify the logical interfaces and not the physical interface on which they reside.

[Network Interfaces]

- **Increase in flow-tap capability**—On the flow-tap application, you can now install a maximum of 100 filters and achieve 100 Kpps throughput. Previously, the limits were 20 filters and 25 Kpps throughput.
- **Restriction on compatibility-mode adtran and verilink**—On 2-port and 4-port channelized DS3 (T3) IQ interfaces, you cannot configure compatibility-mode adtran, or verilink at the [edit interfaces interface-name t3-options] hierarchy level. If configured, the default mode is applied on both the interfaces, that is, no subrating.

[Network Interfaces]

- **Non-support for multiple service sets on a single interface**—When you include dynamic application awareness for JUNOS functionality in a service set, including

IDP profiles, application identification rules, application-aware access list rules, and policy-decision statistics profiles, you can apply only one service set to a single interface.

- **Non-support for connectivity fault management with circuit cross-connect**—M7i routers and M10i routers with the Enhanced Compact Forwarding Engine Board (CFEB-E) do not support connectivity fault management (CFM) with circuit cross-connect (CCC) encapsulation.

[Network Interfaces]

MPLS Applications

- **Make-before-break with ultimate hop popping enabled at an egress router**—Now, when you configure a minimum bandwidth requirement for a new LSP, if there is insufficient bandwidth at the egress router on any of the available VT interfaces, the LSP is not allowed to come up. Previously, the LSP would come up, but with penultimate hop popping enabled. Using penultimate hop popping was inefficient especially at a P2MP bud node (a node that acts as both a transit and egress for different sub-LSPs of the same P2MP LSP) since the upstream node would now have to send two copies of the packet over the same link.

[MPLS]

Platform and Infrastructure

- **Change to show route forwarding-table and show route commands**—The `show route` command no longer displays private, that is, internal, routing tables, and the `show route forwarding-table` command no longer displays private, or internal, forwarding tables. To display output for all routing tables, including private routing tables, use the new `show route all` command. To display output only for private routing tables, use the new `show route private` command. Alternatively, you can use the `show route table routing-table-name` command to display output for a specific private routing table, where *routing-table-name* can be one of the private routing tables, such as `juniper_private1.inet.0`. To display output for all forwarding tables, including private forwarding tables, use the new `show route forwarding-table all` command. To display output for a specific private forwarding table, use the `show route forwarding-table forwarding-table-name` command.

[Routing Protocols and Policies Command Reference]

Routing Protocols

- **Support to specify AS loops value for BGP neighbors per address family**—You can now configure the `loops number` statement for a BGP neighbor for any address family indicator (AFI), or protocol family, supported by BGP. Previously, the `loops` statement was supported only for the global autonomous system (AS) number or the local AS number. When you configure the `loops` statement for a BGP neighbor, the JUNOS Software uses that value to evaluate AS paths received from a BGP peer for the specified address family rather than the `loops` value configured for the AS number at the `[edit routing-options]` hierarchy level.

To specify the maximum number of times a local-AS can appear in an AS path received from a BGP peer for a specific address family, include the `loops number` statement at the `[edit protocols bgp group group-name neighbor address family family-name (any | flow | labeled-unicast | multicast | unicast | signaling)]` hierarchy level. For *number*, specify a value from 1 through 10. You can also configure the `loops` statement at the global and group BGP hierarchy levels for any protocol family supported by BGP.

[Routing Protocols]

- **Traffic engineering enhancement for IS-IS and OSPF**—By default, the JUNOS Software prefers IS-IS routes in the traffic engineering database (TED) over other IGP routes even if the routes of another IGP are configured with a lower, that is, more preferred, preference value. The traffic engineering database assigns a credibility value to each IGP and prefers the routes of the IGP with the highest credibility value. You can now configure IS-IS and OSPF to take protocol preference into account to determine the traffic engineering database credibility value. When protocol preference is used to determine the credibility value, IS-IS routes are not automatically preferred by the traffic engineering database, depending on your configuration. For example, OSPF routes have a default preference value of 10, while IS-IS Level 1 routes have a default preference value of 15. When protocol preference is enabled, the credibility value is determined by deducting the protocol preference value from a base value of 512. Using default protocol preference values, OSPF has a credibility value of 502, while IS-IS has a credibility value of 497. Because the traffic engineering database prefers IGP routes with the highest credibility value, OSPF routes are now preferred.

To configure OSPF to take protocol preference into account to determine the traffic engineering database credibility, include the **credibility-protocol-preference** statement at the `[edit protocols ospf traffic-engineering]` hierarchy level. To configure IS-IS to take protocol preference into account to determine the traffic engineering database credibility, include the **credibility-protocol-preference** statement at the `[edit protocols isis traffic-engineering]` hierarchy level.



NOTE: OSPFv3 is not supported for this feature.

Also, the **show ted protocol** command has been enhanced to display the correct credibility value when protocol preference is enabled.

[Routing Protocols, Routing Protocols and Policies Command Reference]

- **Support for IPv4 and IPv6 routing tables in the same import routing table group**—You can now configure an import routing table group that includes both IPv4 and IPv6 routing tables. Previously, an import routing table group could include only either IPv4 or IPv6 routing tables. To configure an import routing table group, include the **import-rib** *[routing-table-names]* statement at the `[edit routing-options rib-groups group-name]` hierarchy level. For *routing-table-names*, include the names of the routing tables to which you want to import routes.

The ability to configure an import routing table with both IPv4 and IPv6 routing tables enables you, for example, to populate the **inet6.3** routing table with IPv6 addresses that are compatible with IPv4. Specify **inet.0** as the primary routing table, and specify **inet6.3** as a secondary routing table. You must also apply the routing table group to a specific protocol, such as OSPF, by including the **rib-group** *rib-group-name* statement at the `[edit protocols protocol-name]` hierarchy level. Additionally, any configured export routing table group must include IPv4 routing tables only.

[Routing Protocols]

- **Modifying BGP to prepend a local AS number to an AS path during route export**—To modify BGP to prepend only the local AS number to an AS path

during route export, include the `no-prepend-global-as` statement at the `[edit protocols bgp group group-name neighbor neighbor-address local-as autonomous-system-number]` hierarchy level. When you include the `no-prepend-global-as` statement as part of the `local-as` statement at the global, group, or neighbor `[edit protocols bgp]` hierarchy levels, BGP does not prepend the AS number of the master instance to the AS path during route export. Instead, BGP prepends the local AS number to form the session with the peer.

[Routing Protocols]

- **Class E addresses**—The JUNOS Software now allows Class E addresses to be configured on interfaces. To allow Class E addresses to be configured on interfaces, remove the Class E prefix from the list of martian addresses by including the `[edit routing-options martians 240/4 orlonger allow]` configuration statement.
- **Increase in limit to external paths accepted for BGP route target filtering**—You can now specify for BGP to accept up to 256 external paths for route target filtering. Previously, the maximum number that you could configure was 16. The default value remains one (1). To specify the maximum number of external paths for BGP to accept for route target filtering, include the `external-paths number` statement at the `[edit protocols bgp family route-target]` hierarchy level. This statement is also supported for BGP groups and neighbors.

[Routing Protocols]

Services Applications

- **Border Gateway Function (BGF)**—Non-cumulative RECRTCP statistics are properly aggregated at the termination level through use of a weighted average.
- **Border Gateway Function (BGF)**—On M120, M3240, and T640 routers, you might encounter problems implementing a new value for maximum concurrent calls on a virtual BGF when the BGF is in service. Documentation for updating maximum concurrent calls will be corrected to indicate that changing this value requires taking the BGF out of service, configuring the new value, putting the BGF back in service, and restarting the `pgcpd` process.

[Multiplay Solutions, Services Interfaces]

- **New configuration to avoid IDP traffic loss (MX Series routers)**—When the MultiServices DPC configured for a service set is either administratively taken offline or undergoes a failure, all traffic entering the configured interface with an IDP service set is dropped without notification. To avoid this traffic loss, include the `bypass-traffic-on-pic-failure` statement at the `[edit services service-set service-set-name service-set-options]` hierarchy level and (for TCP traffic only) the `ignore-errors tcp` statement at the `[edit interfaces interface-name services-options]` hierarchy level. When you configure these statements, the affected packets are forwarded, in the event of a MultiServices DPC failure or offlining, as though interface-style services are not configured. This issue applies only to MultiServices DPCs on MX Series routers and does not affect MS-400 PICs on M120 or M320 routers.

[Services Interfaces]

Subscriber Access Management

- **Support for classifiers and rewrite-rules with a subscriber interface in dynamic profiles**—You can now associate classifiers and rewrite-rules with a subscriber interface in a dynamic profile. In the current release, this feature is available for testing purposes only.

You must statically configure the classifiers and rewrite-rules at the [edit class-of-service] hierarchy level. To associate a classifier configuration with a subscriber interface in a dynamic profile, include the `classifiers` statement at the [edit dynamic profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

To associate a rewrite-rule configuration with a subscriber interface in a dynamic profile, include the `rewrite-rules` statement at the [edit dynamic profiles *profile-name* class-of-service interfaces *interface-name* unit *logical-unit-number*] hierarchy level.

The *JUNOS Subscriber Access Configuration Guide* does not contain configuration information for this feature.

[Class of Service]

- **Support for statically configured schedulers in dynamic profiles for subscriber access (M120, M320, and MX Series routers)**—You can now configure a combination of static and dynamic parameters for individual schedulers in a dynamic profile for M120, M320, and MX Series routers. In earlier releases, a dynamic profile supported one definition for a dynamic scheduler, which contained scheduler parameters specified using predefined variables. For example:

```
schedulers {
  $junos-cos-scheduler {
    transmit-rate percent $junos-cos-scheduler-tx;
    buffer-size percent $junos-cos-scheduler-bs;
    priority $junos-cos-scheduler-pri;
    drop-profile-map loss-priority low protocol any drop-profile
      $junos-cos-scheduler-dropfile-low;
    drop-profile-map loss-priority high protocol any drop-profile
      $junos-cos-scheduler-dropfile-high;
    drop-profile-map loss-priority medium-low protocol any drop-profile
      $junos-cos-scheduler-dropfile-medium-low;
    drop-profile-map loss-priority medium-high protocol any drop-profile
      $junos-cos-scheduler-dropfile-medium-high;
    drop-profile-map loss-priority any protocol any drop-profile
      $junos-cos-scheduler-dropfile-any;
  }
}
```

Within a dynamic profile, you can choose to configure one dynamic scheduler definition, or combine static and dynamic scheduler parameters in many static scheduler definitions. Combining static and dynamic scheduler parameters enables you to provide subscribers with unique rate configurations that the RADIUS definitions for predefined variables do not allow.

To configure a static scheduler that contains both static and dynamic parameters, include the `schedulers scheduler-name` statement at the [edit dynamic profiles

profile-name class-of-service] hierarchy level. Schedulers that combine static and dynamic parameters must have a specific scheduler name, not the `$junos-cos-scheduler` variable.

In the following example, the network administrator configures the transmission rate for the data service with the `transmit-rate` statement. By specifying the `$junos-cos-scheduler-tx` variable, RADIUS returns the actual percentage value for the transmission rate when the subscriber logs in. The network administrator also specifies the `rate-limit` statement, which limits the transmission rate to the rate-controlled amount during congestion.

For the best-effort service, the network administrator assigns the remaining transmission rate that is available using the `remainder` statement.

```
schedulers {
  data-scheduler {
    transmit-rate percent rate-limit $junos-cos-scheduler-tx;
    buffer-size percent $junos-cos-scheduler-bs;
    priority $junos-cos-scheduler-pri;
    drop-profile-map loss-priority low protocol any drop-profile d0;
    drop-profile-map loss-priority medium-low protocol any drop-profile d1;
    drop-profile-map loss-priority medium-high protocol any drop-profile d2;
    drop-profile-map loss-priority high protocol any drop-profile d3;
    drop-profile-map loss-priority any protocol any drop-profile all;
  }
  best-effort-scheduler {
    transmit-rate remainder
    buffer-size percent $junos-cos-scheduler-bs;
    priority medium-high;
    drop-profile-map loss-priority low protocol any drop-profile
      $junos-cos-scheduler-dropfile-low;
    drop-profile-map loss-priority medium-low protocol any drop-profile d1;
    drop-profile-map loss-priority medium-high protocol any drop-profile
      $junos-cos-scheduler-dropfile-medium-high;
    drop-profile-map loss-priority high protocol any drop-profile d3;
    drop-profile-map loss-priority any protocol any drop-profile
      $junos-cos-scheduler-dropfile-any;
  }
}
```

[Subscriber Access]

- **RADIUS accounting enhancements**—Configures the router to request that the RADIUS accounting server immediately update accounting statistics when a change of authorization (CoA) occurs. You use the `coa-immediate-update` statement at the [edit access profile *profile-name* accounting] hierarchy level to configure the accounting update.

The following Juniper Networks VSAs (vendor ID 4874) are now included in the RADIUS accounting statistics.

- Ingress-Policy-Name—VSA 26–10
- Egress-Policy-Name—VSA 26–11
- DHCP-Options—VSA 26–55

VSA 26-10 and 26-11 are specified by the `$junos-input-filter` and `$junos-output-filter` predefined variables, which are used in dynamic profiles. See “Dynamically Attaching Filters Using RADIUS Variables” in the *JUNOS Subscriber Access Configuration Guide*.

[Subscriber Access]

- **Change of authorization (CoA) message enhancement**—When you change class-of-service (CoS) services for a single dynamic subscriber interface using a RADIUS change of authorization (CoA) message, scheduler map names are now replaced with the value in the CoA message. The behavior for single subscribers is now the same for multiple subscribers enabled on the logical interface with the `aggregate-clients replace` statement.

[Subscriber Access]

- **Enabling and disabling DHCP snooping support**—You can now explicitly enable or disable DHCP snooping support on the router. If you disable DHCP snooping support, the router drops snooped DHCP discover and request messages.

To enable DHCP snooping support, include the `allow-snooped-clients` statement at the [edit forwarding-options dhcp-relay overrides] hierarchy level. To disable DHCP snooping support, include the `no-allow-snooped-clients` statement at the [edit forwarding-options dhcp-relay overrides] hierarchy level. Both statements are also supported at the named group level and per-interface level.

In JUNOS Release 10.0 and earlier, DHCP snooping is enabled by default. In release 10.1 and later, DHCP snooping is disabled by default.

[Subscriber Access]

- **RADIUS VSA 26-56 (vendor ID 4874) format**—The RADIUS DHCP-MAC-Address attribute (VSA 26-56) contains the subscriber's MAC address when available in RADIUS Authentication, Accounting-Start, Interim-Accounting, and Accounting-Stop requests. VSA 26-56 now appears in an ASCII string in hexadecimal format. In previous releases, the ASCII string appeared as a 6-byte binary value.

[Subscriber Access]

- **RADIUS interim accounting**—When subscriber management receives the RADIUS Acct-Interim-Interval attribute (attribute 85), RADIUS interim accounting is performed based on the value in the attribute. The router uses the following guidelines:

- Attribute value is within the acceptable range (10 to 1440 minutes)—Accounting is updated at the specified interval.
- Attribute value of 0—No RADIUS accounting is performed.
- Attribute value is less than the minimum acceptable value (10 minutes)—Accounting is updated at the minimum interval.
- Attribute value is greater than the maximum acceptable value (1440 minutes)—Accounting is updated at the maximum interval.

In previous releases, a RADIUS attribute set to zero (0) prevented subscribers from connecting.

[Subscriber Access]

User Interface and Configuration

- Beginning with JUNOS Release 9.6, the **show system switchover** command displays the following error message when you run the command on the master Routing Engine of a router: “error: the kernel-replication subsystem is not running.”. This is because the kernel replication process does not run on the master Routing Engine. This process runs only on the backup Routing Engine.

VPNs

- **Regular expressions for VRF import policies**—You can now include regular expressions as a part of VRF import policies. For example, you can configure the following using the **community** statement at the **[edit policy-options policy-statement *policy-statement-name*]** hierarchy level: **community HIGH_PRIORITY members *:50**.



NOTE: You cannot configure a regular expression as a part of a route target extended community.

[VPNs]

- Related Topics**
- New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 6
 - Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 54
 - Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers on page 96
 - Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 103

Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers

The current software release is Release 9.6R4. For information about obtaining the software packages, see “Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers” on page 103.

- Current Software Release on page 55
- Previous Releases on page 76

Current Software Release

Outstanding Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers

Class of Service

- On MX Series routers with Enhanced DPCs, bandwidth sharing between two schedulers, one with high and the other with strict-high priority, might not be as expected when the schedulers are oversubscribed. That is, only one queue can use all of the excess bandwidth. This issue occurs when the schedulers are configured on logical interfaces. [PR/265603]
- When you set the port speed of a multirate SONET Type 2 PIC to OC3, the CoS speed value is not changed correctly within the Packet Forwarding Engine. The speed value remains OC12, which results in unexpected CoS behavior. There is no workaround. [PR/279617]
- Configuring rewrite-rules on PPPoE IFLs is not supported. [PR/438327]
- In the cosd logs, "entries" is misspelled as "enteries." [PR/439993]
- On M Series and T Series routers, the forwarding class information is lost when the packet enters the GRE tunnel with a clear-dont-fragment-bit enabled. Additionally, on an Enhanced FPC or M120 FEB, the packet is also likely to be dropped if it is classified to a packet loss priority (PLP) other than low. [PR/514162]

Forwarding and Sampling

- On M320 and T Series routers, 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]
- While the JUNOS Software adopts random as its sampling algorithm, the SAMPLING_ALGORITHM in the jflowv9 template shows 0x01 (deterministic) instead of 0x02 (random). [PR/438621]
- Under rare circumstances, if the filter is changed while a counter query is in progress and the system is under heavy load, the system might crash. [PR/447033]
- The numerical values configured for the ip-options match criteria on a firewall filter match any ip-options no matter what is specified. [PR/516778]
- On T640 and T1600 routers with ST chipset FPCs, in some cases when the IPv6 firewall filters with match conditions configured on address prefixes are longer than 64 bits, the filter might not be evaluated correctly. This can lead to loss of packets. [PR/524809]

High Availability

- The primary Routing Engine might lose the CM/CP information if it loses connectivity with the redundant Routing Engine (i.e., by disabling GRES, or halting and rebooting the redundant Routing Engine). This can cause small packet drop on multicast traffic upon a multicast distribution tree change. [PR/278882]
- A problem occurs during graceful Routing Engine switchover (GRES) when a static route pointing to a private interface such as fxp0 is created using the passive retain option. It is recommended to not use the **passive** option along with the static route on the private interface. [PR/412746]
- When a Routing Engine switchover occurs at the same time that FRUs are re-connecting to the Routing Engine, kernel panic can occur. [PR/419966]
- The SSH keys are not in sync between the master and backup Routing Engine when SSH is enabled after a graceful Routing Engine switchover (GRES). [PR/455062]

Interfaces and Chassis

- 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]
- 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]
- When the ATM scheduler map is programmed, the code does not check if the early packet discard (EPD) configured on the forwarding class exceeds the max_epd that the hardware supports. [PR/70336]
- 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]
- Hot swapping the M120 router fan tray might cause the **Check CB** alarm to activate. [PR/268735]
- On the JCS1200 platform, 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 using 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 an MS 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]

- The following messages are displayed on both the primary and secondary RLSQ MS 500 PICs: "SCHED: %PFE-0: Thread 7 ran for x ms without yielding," "Scheduler Oinker." [PR/286357]
- 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 (GRES). [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]
- When forwarding-options is configured without route-accounting, the commit completes with the message, "Could not retrieve the route-accounting." However, no functionality is affected. [PR/312933]
- The backup Routing Engine can fail to obtain mastership in the following cases:
 - Re0 gets stuck and doesn't reboot.
 - Due to some hardware problem, re0 loses its connectivity with both the Control Board and the Packet Forwarding Engine.

[PR/405412]

- 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 a backup Routing Engine is replaced after a graceful Routing Engine switchover (GRES), the device control process (dcd) generates a new link local address on non-MAC interfaces such as SONET. [PR/429078]
- When an IRB and a AE (LAG) interface with two or more child links are in the downstream interface list of a multicast group, traffic forwarding might not work correctly over the IRB. [PR/437436]
- When a PIC is hot swapped on an MX FPC, it could cause the chassisd to restart several times when the PIC is brought online. To prevent this issue, insert the PIC while the FPC is offline. [PR/438971]
- When you configure the **payload port-data** statement at the [edit family mpls hash-key] hierarchy level on M120, M320, or MX Series routers with E3 FPCs, the hashing algorithm might not take the port data values into account. [PR/442223]
- When configured for WAN-PHY framing, the ports on the 4-port 10-Gigabit Ethernet PIC always report zero for path-level errors (BIP-B3) in the output of the **show interfaces extensive** command.

After the fix, the BIP-B3 counter increments when path-level errors occur. However, this counter is an approximation and not an accurate accounting of the path-level errors that actually occur on the link. [PR/447653]

- The 10GE XENPAK interface might flap when the transmission gear fails over. [PR/446973]
- M7i routers and M10i routers with Enhanced Compact Forwarding Engine Board (CFEB-E) do not support connectivity fault management (CFM) with circuit cross-connect (CCC) encapsulation. [PR/449684]
- If virtual tunnel PICs and ingress traffic manager is enabled on the same Packet Forwarding Engine/PIC on an EQ DPC, then the SNMP walk of the interface can time out. [PR/458565]
- Both the working and protect circuit are stuck in the “disabled” state when the TX cable is unplugged and the RX cable is plugged for protect circuit after an Automatic Protection Switching (APS) switchover. [PR/466649]
- When loopback is configured on t3 under ct3, t1 under ct1, or e1 under ce1, no error syslog message is logged. Additionally, the **show interface extensive** command on the t3/t1/e1 displays "loopback" even though it is not actually applied. [PR/486424]
- On MX Series routers, the traffic is forwarded over the backup link even after the primary link is disabled and enabled again. [PR/493861]
- When **t1-options** are configured at the **[edit interfaces ct1-x/y/z]** hierarchy level, some ct1 interfaces of a 10xCHT1 IQ PIC might flap when the configuration changes are committed. As a workaround, remove the **t1-options**. [PR/500820]
- When the **show lacp interface aex** command is used for a nonexistent AE interface, no error is returned. [PR/503806]
- If a T640-FPC4-ES is installed in a T1600 router and an SIB statistics collection is performed, the message log might report "JBUS: U32 read error, client." only if one of the SIBs is faulted or in the offline state. This system log message will also appear if the T640-FPC4-ES FPC is removed from the chassis. There is no operational impact. [PR/504363]
- Under certain circumstances, the E3 IQ PIC might report bogus CCV, CES, and CSES alarms. [PR/505921]
- Under certain circumstances, the chassisd process might crash on a backup Routing Engine while a configuration is committed. [PR/512044]
- On IQ2 and IQ2E 10GE PICs operating in WAN-PHY mode, the path trace information is not transmitted to the remote end. [PR/518331]

Layer 2 Ethernet Services

- DHCP packets cannot be processed on an auto-sensed VLAN interface if the DHCP configuration for the interface is performed after the auto-sensed VLAN interface is instantiated. As a workaround, clear the auto-sensed VLAN interface(s) after the DHCP configuration is made for the interface(s). [PR/417958]
- While inserting the DPC into the chassis, the chassisd log can display a bogus error message: "FPC X temperature is -60 degrees C, which is outside operating range." This message does not impact any functionality. [PR/470512]

MPLS Applications

- The `rt` column in the output of the `show mpls lsp` command and the active route counter in the output of the `show mpls lsp extensive` command are incorrect when per-packet load balancing is configured. [PR/22376]
- The routing protocol process might crash at `rsvp_find_ip_tag_route` occasionally. [PR/55748]
- No point-to-multipoint LSPs are reported when the `show mpls lsp p2mp` command is issued. As a workaround, execute the `show mpls lsp` command before you execute the `show mpls lsp p2mp` command. [PR/266343]
- For point-to-multipoint label-switched paths configured for VPLS, the `ping mpls` command reports a 100 percent packet loss even though the VPLS connection is active. [PR/287990]
- Constrained Shortest Path First (CSPF) fails to calculate a P2MP LSP reroute path merger upon a user configuration change. [PR/454692]
- When an RSVP LSP is configured with the `no-install-to-address` option and is not associated with CCC connection flaps, the routing protocol process will crash when the LSP comes up again. To avoid the problem, make sure that the LSP is either a transmit LSP for a CCC connection or that the `install` option is also configured on the LSP. [PR/471339]
- The `show route table mpls.0 label-switched-path lsp-name` command might cause the routing protocol process to crash if no route is found. [PR/507239]
- The targeted LDP neighbor might remain up with an old IP address that was previously in use with the loopback on the remote neighbor. This can happen when the following is performed on the remote neighbor:
 - A secondary loopback (lower than the current primary) is added and no primary keyword is associated with either address.
 - A second loopback address is added with the primary keyword.

This will result in the targeted LDP neighbors being up with both the IP addresses. The neighbor with the old address can remain up even after the old loopback address is deleted on the remote neighbor. The neighborship with the old address eventually times out if the router-id is changed to reflect the new loopback address on the remote neighbor. [PR/518102]

Network Management

- The interface description configured on the logical instance displays only on a `commit full` and not with the `commit` command. [PR/288595]
- `Tcpdump` might crash when IPv6 malformed packets are received with `NextHeader = AH`. [PR/399073]
- When the `size` option is not set using the `monitor traffic` command, the host-generated packets (OSPF, LDP, and so on) are truncated. [PR/437655]

- After changes are made to the firewall, and the counters are cleared and committed, SNMP sends the wrong value for 5 seconds and a discrepancy between the cli output and get snmp output occurs. [PR/459583]
- Under certain conditions, the SNMPD crashes due to a BAD_PAGE_FAULT. [PR/496351]

Platform and Infrastructure

- On T Series routers, a Layer 2 maximum transmission unit (MTU) check is not supported for MPLS packets exiting the routing platform. [PR/46238]
- When you configure a source class usage (SCU) name with an integer (for example, 100) and use this source class as a firewall filter match condition, the class identifier might be misinterpreted as an integer, which might cause the filter to disregard the match. [PR/50247]
- 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]
- 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 T Series routers, the commit operation succeeds when you include the `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level, but MPLS labels are still included in the hash key. [PR/80334]
- Traceroute does not work when ICMP tunneling is configured. [PR/94310]
- Initiate doesn't parse the configuration present in `init.conf` regardless of position. [PR/94576]
- 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]
- On T Series and M320 routers, multicast traffic with the "do not fragment" bit set is being dropped due to configuring a low MTU value. The router might stop forwarding all traffic transiting this interface if the `clear pim join` command is executed. [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]

- 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]
- On a Protected System Domain, under the following conditions an FPC might generate a core file and stop operating:
 - A firewall policer with a large number of counters (for example, 20,000) is applied to a shared uplink interface, and
 - 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]

- When a CFEB failover occurs on an M10i or M7i router with 4000 or more IFLs, the following message will display:

```
IFRT: 'IFD ioctl1' (opcode 10) failed
ifd 153; does not exist
IFRT: 'IFD Ether autonegotiation config' (opcode 163) failed
```

The message has no operational impact. When the backup CFEB becomes the active CFEB, the message will not display. [PR/400774]

- The following error message might display for tunnel PICs in **/var/log/messages**: “/kernel: if_tunnel_cookie_remove no callback!!!”. These messages are harmless and are not valid. [PR/422715]
- On M120, M320, MX Series, and T Series routers, traceroute leaving an LSP configured for explicit-null and no-decrement-ttl or no-propagate-ttl, might not show the transit IP hop router immediately after the LSP egress router. [PR/438735]
- On M7i routers, kernel panic might occur during route changes. [PR/439420]
- In some cases, the alarms displayed in FPM and the alarms shown using the **show chassis alarms sfc 0** command do not match. [PR/445895]
- The SFC management interface **em0** is often displayed as **fxp0** in several warning messages. [PR/454074]
- If the subinterface on an aggregate interfaces goes down, the GRE traffic egressing that interface might not use the backup subinterface, resulting in the GRE traffic being dropped. [PR/454751]
- When the strict-high priority queue is overloaded, the high priority queue might starve, resulting in the loss of high-priority traffic. [PR/455152]

- When some DHCP related configurations are added or deleted (for example, `delete bootp server address`), under some rare conditions while incorporating these changes, the router might generate an FUD core. [PR/458132]
- When the flow monitoring version 9 feature is enabled on an MS PIC (or service PIC which supports flow monitoring version 9), the MS PIC might crash upon receiving certain corrupted IPv6 packets. [PR/458361]
- The VPN label does not get pushed on the label stack for Routing Engine-generated traffic with `l3vpn-composite-next-hop` activated. [PR/472707]
- Payload corruption and packet drops might occur for packets larger than 3000 bytes when MPLS over GRE is configured on a service PIC. [PR/478563]
- An invalid IP protocol version is served as a valid version. The JUNOS router forwards IP packets with the version field set to values other than 4 and 6; for example, 11 or any (unassigned). [PR/481071]
- The `fxp0` packet counter statistics are inconsistent between the physical interface and the logical interface as the statistics are updated twice. [PR/486200]
- A problem occurs on an M120 router with an FEB redundancy configuration when the backup FEB is protecting a non-primary FEB. In this case, the Routing Engine will prompt the incorrect Packet Forwarding Engine for status, causing delays in the SNMP responses. [PR/490172]
- If you configure an IP address with a larger subnet, for example, /19, on a different interface first and the router begins to negotiate for the ARP of a specific host on that interface, it can get stuck in a hold state. If you later configure a more specific subnet of /29 on another interface from where the host can be reached, the forwarding table will still prefer the route with the hold entry via /19 instead of the route with the ucst entry via /29. [PR/491468]
- The MAC address of a configured static NDP entry is overwritten when NA is received from a connected device. [PR/499418]
- The static NDP entry remains permanent if the `refcount` is more than 1, even after the static configuration is deleted. [PR/499441]
- The tty sessions to a router can cause a null pointer dereference. [PR/502816]
- When an AE interface on an ECMP path is taken down, packet drops can occur on traffic that is on another link in the ECMP path. [PR/513102]
- Setting the TCP maximum segment size (MSS) might not change the actual MSS value. [PR/514196]
- On M120 and MX Series routers when an AE interface (with LACP enabled) is used as a core-facing interface for L3VPN, the non MPLS traffic received on the AE interface can sometimes get black-holed. To recover from this state, deactivate and reactivate the AE interface in the configuration. [PR/514278]
- When the Destination Class Usage (DCU) is configured with unicast Reverse Path Filter (uRPF) and egress forwarding-table filter within the VRF, a VPN route flap might trigger a jtree memory leak. [PR/521609]

Routing Policy and Firewall Filters

- If a routing protocol running an MSDP receives an SA that is filtered via the MSDP import policy, it will still create a forwarding entry if it subsequently receives a (*,G) join for that group. [PR/63053]

Routing Protocols

- 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 without any defined peers, a warning message no longer 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]
- 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]
- The `keepalive timeout` counter for multicast sessions might not display after you deactivate and activate the `pim` protocol. This is a cosmetic issue and there is no interruption to the multicast traffic flow. [PR/419509]
- Setting the `advertise-high-metric` option while using IS-IS overload also suppresses route leaking. [PR/419624]
- The rendezvous point (RP) is not being learned on a router with auto-rp discovery configured, when there is a mismatched PIM interface configuration on a router with auto-rp discovery configured and on a router with auto-rp mapping configured. For example, one router having an IFL with PIM configured and the other having an IFL with PIM disabled. As a workaround, ensure that PIM is enabled on all IFLs on both routers. [PR/445917]
- On a router with VPNs configured, modifying or adding configuration might cause the "age" of the secondary routes to reset to 0. [PR/447802]
- The backup Routing Engine can generate routing protocol process and kernel cores if BGP damping is configured along with nonstop active routing (NSR). [PR/452217]
- If the routing protocol process (rpd) experiences a restart, it might not receive the first PIM hello packet from a PIM neighbor after the restart. This can delay the establishment of PIM neighbors, and therefore multicast traffic convergence, for up to twice the PIM hello interval. [PR/452751]
- On JUNOS OSPF, all locally generated Type 5 LSAs are purged and regenerated while deleting an NSSA area from the area border router (ABR). [PR/457579]

- The BGP strip confederation logic does not include the number of memory segments to check. This leads to it running on random data, causing the routing protocol process (RPD) to core. [PR/465624]
- When a Flexible PIC Concentrator reboots or an interface is temporarily deactivated, two RPD_PIM_NBRDOWN messages are logged for every PIM neighbor affected. However, only one RPD_PIM_NBRUP message is logged when the service is restored. This could lead to inconsistencies in any management software. [PR/472873]
- When a dampened route is restored, the accepted count for the peer in the **show bgp summary** output does not increment. [PR/473567]
- When a PIC with a PIM-enabled interface is brought online, the router might send the first PIM hello slightly before the interface comes up. This causes the router to drop the first PIM hello message to its neighbor. [PR/482903]
- During transient periods where both a secondary and primary LSP exist in a routing table, and the number of LSP NHs is greater than 16 in a multigateway scenario, IS-IS might remove the preferred LSP NH. For example, IS-IS could remove an HIPRI LSP. [PR/485748]
- The routing protocol process crashes at task_reconfigure in task.c:2653 during a failed MVPN configuration change. [PR/486183]
- The Routing Protocol does not process the PIM register messages from a first-hop router in an IPv6 embedded routing protocol group when the Register message does not have the NULL bit set. [PR/486902]
- On receiving a BGP open message with a hold time of 0 seconds, the JUNOS Software ignores that value and sets the hold time to 90 seconds. [PR/487107]
- The BGP BMP message for IPv6 withdraw encoding does not follow the BMP-draft. [PR/512780]
- In route reflector and ASBR VPN scenarios, the routing protocol process might crash when changes occur to a prefix in the primary table at the same time as BGP tries to send out updates via the secondary table. [PR/515626]
- The configured robust count value is not applied on the non-querier router when it receives a robust count value of 0. It uses the default value (2) instead of the configured value. [PR/520252]
- The tag_encoder is unable to handle attempts to stack EXPLICIT_V6_ NULL (label 2) over an existing stack with label2 on top. Additionally, the BGP module does not send label 2 when readvertising a prefix from an inet6 unicast session to a inet6 labeled-unicast session. [PR/523824]

Services Applications

- The **show services accounting flow-detail extensive** command sometimes displays incorrect information about input and output interfaces. [PR/40446]
- When a routing platform is configured for graceful Routing Engine switchover (GRES) 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]

- Detection of failure of remote PPP clients on the LNS through LCP echo requests will take longer due to the increase in the number of echo request retries. [PR/250640]
- When using L2TP services on M Series routers, every session or tunnel connection and disconnection causes memory leak. [PR/312961]
- With the E-CFEB on M7i and M10i routers, If you configure a firewall filter with an action of sampling and then apply the filter to the interface, all packets received on the PIC are corrupt and consequently dropped. [PR/408802]
- Flow monitoring records are not generated as fragmented IPv6 packets are not getting sampled. [PR/478571]
- When the Border Signaling Gateway (BSG) configuration contains a policy that has a term with regular expressions, configuration changes do not take effect immediately after you receive the message that the commit process is complete. The time it takes for the configuration to take effect depends on how many regular expressions are in your term.

For example, if you have a term with four regular expressions, configuration changes do not take effect until 50 seconds after you receive the message that the commit process is complete. This behavior occurs whether you have a list or regular expressions (for example, regular-expression [sip:88824.* sip:88821.* sip:88822.sip:88823.*]) or you group regular expressions using the | symbol (for example, "sip:88821.*|sip:88822.*|sip:88823.*|sip:88824.*").

The time that it takes the software to apply the configuration changes increases exponentially with the number of regular expressions in your configuration. [PR/448474]

- Configuring different autonomous-system-types (origin and peer) towards two v5 servers does not work and origin is taken as the autonomous-system-type for both flow servers. [PR/496954]
- When a backup gateway is configured in any term under IPsec stanza, for any subsequent terms where this backup gateway is configured as the primary, IPsec tunnel establishment will fail. [PR/510608]

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]
- The router always uses the **revert-interval** value that is configured at the [edit access] hierarchy level, and ignores any **revert-interval** value that is configured at the [edit access profile] hierarchy level. If no value is configured, the router uses the default value of 600 seconds. [PR/454040]
- The DHCP clients might not get bound after a filter action under a firewall filter context is deactivated and deleted. [PR/488627]

User Interface and Configuration

- The CLI does not warn if multiple users are configured with the same user-id. [PR/55774]
- The user cannot prevent the deletion of configuration groups with the `allow-configuration` and `deny-configuration` statements. [PR/59187]
- When the `get-configuration` or `load-configuration` commands are run using JUNOScript, these events are not recorded in the syslog. [PR/64544]
- 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 JUNOScript perl module for netconf does not support configuration-text. [PR/82004]
- The “Local Password:” is prompted even though the authentication order has a password configured. [PR/94671]
- 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]
- The “replace:” tag is missing from the output of the `save terminal` command from inside a configuration object.

Example:

```
edit system
save terminal
system {
    host-name blue;
}
```

[PR/269736]

- The user can still commit an invalid configuration successfully, even when DDL checks exist. [PR/282896]
- After AI scripts are added, the existing management sessions (including the one used to add the AI scripts) must exit the `edit` mode and reenter for any subsequent configuration changes to take effect. Changes made in these existing `edit` sessions are not written to the candidate configuration. [PR/297475]
- A user class configuration with the deny command `".*"` returns `.noop` error when `enter` is used on the router CLI. As a workaround, replace `"^$"` with `"^.noop-command$"` in allow regex. [PR/311426]
- Users who have superuser privileges will sometimes have their access restricted to view permission only. [PR/388053]
- When the filter `config-text` is used in the NETCONF `get-config` command, a syntax error occurs and the router configuration cannot be returned in ASCII format. [PR/430799]

- On M Series, MX Series, and T Series routers, the user cannot differentiate between active and inactive configurations for system identity, management access, user management, and date and time pages. [PR/433353]
- Selecting the Monitor port for any port in the Chassis Viewer page takes the user to the common Port Mirroring page instead of the corresponding Monitoring page of the selected port. [PR/446890]
- The core files cannot be removed using the **file delete** command unless the Routing Engine name is included in the path. [PR/469168]
- If the time zone is set to 'Europe/Berlin', the command **commit** fails at "time-string." [PR/483273]
- On M7i and M10i routers with Enhanced CFEB installed, the chassis viewer plugin does not display the Routing Engine in the front view and the E-CFEB in the rear view. However, the chassis contents from the system (left side tab) displays all the list of components correctly. [PR/483375]
- If the user in the Backup Routing Engine with config-private mode activates graceful Routing Engine switchover (GRES) and uses commit synchronize, a synchronization error might occur during GRES switchover. [PR/486637]
- On J-Web, the error message: "Fatal error: Allowed memory size..." displays when the Interfaces tab is selected. This message also displays when the Interfaces tab under Class-of-Service is selected. [PR/495825]
- The **load replace** command does not consider the allow-configuration configuration. [PR/501992]

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]
- Once a VPLS NLRIS is received with the Local Preference value of zero, it is assumed that the remote site is not designated even if there is one remote site. As a workaround, use a non-zero Local Preference value. [PR/70601]
- 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]
- Under certain circumstances, if BGP is configured as the PE router to CE router protocol in a Layer 3 VPN routing instance, renaming the routing instance can cause the PE router to CE router session to stay down. [PR/399275]
- The IPv6 multicast packet forwarding fails when a VT interface is configured for multicast in the egress PE with NGEN-MVPN. [PR/431957]
- On MX, M120 and new EIII FPCs on M320 routers, the ISO/Connectionless Network Service (CLNS) packets over the translational cross-connect (TCC) are

dropped in the case of frame relay, even though the family TCC has been configured to switch family iso on the frame relay interface. [PR/462052]

Resolved Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers

Class of Service

- On the Qchip, the shaping accuracy is affected by the configured logical interface shaping rate. [PR/79319: This issue has been resolved.]
- DHCP traffic might stop being processed for some subscribers under heavy login and logout conditions when the 802.1 classifiers are in use. [PR/470513: This issue has been resolved.]
- The following operations might result in large incorrect queue statistics on IQ2 interfaces:
 - When the IQ2 PIC is restarted, or the interface is deactivated and reactivated, while traffic is on and the configuration defines a high priority queue on the interface.
 - When the high priority queue number is changed under the class-of-service configuration while traffic is on.

[PR/489049: This issue has been resolved.]

- On M Series (except M120 and M320) routers, packet classification will not work on aggregated Ethernet bundles that have LACP enabled. [PR/492057: This issue has been resolved.]
- The class-of-service process crashes on commit if a scheduler-map definition does not have any forwarding-class statement. [PR/499755: This issue has been resolved.]
- On an Ichip-based platform for strict high priority queue (SHQ), the buffer size allocated by the Packet Forwarding Engine is capped by the tx-rate. If the tx-rate is configured to a very small value or is not configured, and is automatically allotted a zero or a very small remaining value; the queue is also allotted a proportionately small delay buffer. This can sometimes lead to Red and Tail drops on the SHQ when there is a burst of traffic (with a certain traffic pattern) on it. As a workaround, configure a nominal tx-rate value (5 percent) for the SHQ. [PR/509513: This issue has been resolved.]

Forwarding and Sampling

- Policers cannot be modified after a system upgrade due to a flaw in the parser routine. This error occurs when the current item is deleted and 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: This issue has been resolved.]
- Under certain conditions for prefix optimization, the firewall compiler might discard a prefix configured for accept. This issue depends on the set of prefixes configured to match across the various terms. [PR/486633: This issue has been resolved.]

- A JUNOS Software compiler bug in the match combination optimization could cause an incorrect firewall filter evaluation. [PR/493356: This issue has been resolved.]
- When the MS PIC used for an RLSQ interface resides on an E3 FPC (M320), traffic might stop flowing across the RLSQ interface after the policer on the interface is deactivated. [PR/498069: This issue has been resolved.]
- When a Layer 2 policer is configured under a logical interface having multiple families configured under it, and the policer is changed to another, the newly configured policer might not take effect unless the policer configuration is deactivated and activated. [PR/501726: This issue has been resolved.]
- When a filter group is configured on an interface residing on an ES FPC, the rpf-check configured on that interface will not function correctly. As a workaround, deactivate the configured filter group. [PR/503609: This issue has been resolved.]
- After configuring a three-color-policer, a dfwc core file is generated. [PR/509742: This issue has been resolved.]

High Availability

- On an ISSU upgrade from JUNOS Release 9.3 to any of the current higher releases, the ATM logical interfaces will flap. [PR/491511: This issue has been resolved.]

Interfaces and Chassis

- Configuration of duplicate virtual-ip addresses is not allowed across routing instances and logical systems. [PR/402235: This issue has been resolved.]
- On M Series and MX Series routers, the ifHCInOctets retrieved by SNMP might report an incorrect value. [PR/420985: This issue has been resolved.]
- For an IQ2 PIC logical interface, the Input Bytes counter and the Input Packets counter might occasionally be incorrect. The statistics are incorrect when there is significant local traffic associated with the logical interface:
 - The transit Input Bytes and Packets counters for a short duration might count backwards or reset to zero.
 - The Total Input Bytes and Packet counters for a short duration might count backwards.

This issue is transient and happens only during steady traffic flow with significant local traffic. If the traffic is stopped or if the local traffic is marginal compared to the total traffic for the logical interface, then the counters will become accurate. [PR/422109: This issue has been resolved.]

- Under some conditions, if an interface flaps for an interval less than the hold-down time configured value, that interface might stop forwarding even though it shows as being UP. As a workaround, enable traffic monitoring on the interface, or enable and disable the interface. [PR/423065: This issue has been resolved.]
- CFMD might crash when the following are configured and committed at once on a VPLS setup:

- Encapsulation VLAN-VPLS on a physical and logical interface
- Family VPLS on a logical unit
- Interface is added in the VPLS routing instance

As a workaround, add the above configurations one at a time and commit.
[PR/440108: This issue has been resolved.]

- The `show interfaces diagnostics optics` command displays wrong diagnostic information for the SumitomoElectric SFP with vendor part number SCP6F44-J3-ANE. [PR/463837: This issue has been resolved.]
- SFPs are absent in the `show chassis hardware` output following TOXIC SFP messages. [PR/480828: This issue has been resolved.]
- In some cases during the periodic error status monitoring, error messages such as “Wi seg ucode discards in fabric stream” can be displayed on adjacent streams. These messages are cosmetic and can be ignored. [PR/481344: This issue has been resolved.]
- On a TX Matrix router, commit returns a validation error if there are no `fxp0` configurations in the `[groups lccX]` hierarchy level, and the following is applied simultaneously:

```
groups {
  int-disable {
    interfaces <*> disable
    interfaces {
      <*> {
        disable;
      }
    }
  }
}
```

[PR/482612: This issue has been resolved.]

- During graceful Routing Engine switchover (GRES), if the peer's discovery state is passive, the LFM state machine should be kick started even if the kernel state is `SEND_ANY`, otherwise the peer will be stuck in `PASSIVE_WAIT` state. As a workaround, configure both sides in the link-discovery mode as “active.” [PR/490886: This issue has been resolved.]
- The DPC remains in the ready state and the demux0 interface remains in a down state after a chassisd restart without graceful Routing Engine switchover (GRES) enabled. [PR/492961: This issue has been resolved.]
- When an SCB which has an active plane is powered down, HSL link error occurs on unrelated SCBs. [PR/493151: This issue has been resolved.]
- The AE logical interface flaps when the PIC that has the active link-protection member link is taken offline. [PR/493492: This issue has been resolved.]
- The CLI does not respond when Control + c is entered at the “more” separator. [PR/493881: This issue has been resolved.]
- The system might generate a core file when the DPC is removed before it is taken offline. [PR/494625: This issue has been resolved.]

- An outer virtual LAN tag is not added in a provider edge-customer edge link when VPLS traffic arrives with an MPLS value of 2, 3, 4, or 5. However, VPLS traffic with a value of 0, 1, 6, or 7 does not have this issue. [PR/495555: This issue has been resolved.]
- The one-port OC12-3 PIC cannot support eight queues when the **no-concatenate** option is configured. [PR/499452: This issue has been resolved.]
- Polling **iflnOctets** on Gigabit Ethernet IQ PIC VLANs might momentarily return a higher value. [PR/500852: This issue has been resolved.]
- During a link UP/DOWN transition, **jsscd** might crash as a result of a NULL message dereferencing by **jsscd**. [PR/502745: This issue has been resolved.]
- Occasionally, a backup Routing Engine reboot followed by a Routing Engine failover can cause LACP to flap, causing ae bundles to flap. [PR/502937: This issue has been resolved.]
- When an ATM AIS cell is received from the virtual channel under **vlan-vci-ccc** encapsulation, the logical interface will be incorrectly marked as down. There is no workaround. [PR/503653: This issue has been resolved.]
- When **native-vlan-id** is configured for aggregated interface with the child links on an IQ2 PIC, the LACPs are dropped and the links go down. [PR/507040: This issue has been resolved.]
- The **show interfaces diagnostics optics interface** command does not display the unit of measurement when the received power is in a very low range (power < 5e-10). It shows the value of 0.00 without any unit of measurement. [PR/507653: This issue has been resolved.]
- When the master Routing Engine is down and the backup Routing Engine is rebooted, the backup Routing Engine reboots as backup. It does not become the master for five to six minutes. [PR/507724: This issue has been resolved.]
- On MX Series routers, the **chassisd** crashes when the SCB is taken offline and removed. [PR/510950: This issue has been resolved.]
- On M7i and M10i routers, the **syncer** process writes to the file **/var/rundb/chassisd.dynamic.db** every 30 seconds. [PR/511901: This issue has been resolved.]
- Due to a flaw in implementation, the execution of the **show interfaces mac-database** command causes the IQ2 PIC to reboot with the core. [PR/513407: This issue has been resolved.]
- The output of the **show chassis hardware** command might not display the SIB details when the SIB is inserted in the slot. [PR/515789: This issue has been resolved.]
- On some XENPAK modules, the output of the **show chassis hardware** command shows the message "NON-JNPR UNKNOWN" when the FPC is booted. There is no impact on the traffic. To solve this issue, take the PIC offline and bring it back online. [PR/516411: This issue has been resolved.]

MPLS Applications

- Sometimes, a traffic engineered label-switched path that is down does not get re-signaled. [PR/478375: This issue has been resolved.]
- The NGEN-MVPN multicast traffic might be dropped at the ingress router if a point-to-multipoint LSP reoptimization is performed. [PR/491533: This issue has been resolved.]
- A rare condition between the MVPN and RSVP P2MP signaling leads to the creation of stale flood next hops. [PR/491586: This issue has been resolved.]
- An incorrectly changed LDP session authentication key causes the LDP session to fail, and the LDP/IGP synchronization feature stops working. The IGP continues to advertise the link at normal metric values. [PR/499226: This issue has been resolved.]
- LDP might not handle certain error conditions gracefully when NSR is enabled. This might cause the LDP replication state to be stuck in the "In Progress" state forever. [PR/505043: This issue has been resolved.]

Network Management

- Under certain SNMP conditions, the following log message is displayed:

```
M10i-RE0 pfed: PFED_NOTIF_GLOBAL_STAT_UNKNOWN: Unknown global
notification stat: transit options/ttl-exceeded (re-injected)
M10i-RE0 pfed: PFED_NOTIF_STAT_UNKNOWN: Unknown notification type stat:
Unknown
```

This log message might also be displayed during the installation of AI Scripts (version 2.1R2 or above) on the router. AI Scripts versions prior to 2.1R2 do not cause these messages. This is a cosmetic message, and does not have any impact. [PR/427590: This issue has been resolved.]

- When `monitor traffic matching x` is used on RLSQ bundles, no outbound packets are displayed. [PR/468959: This issue has been resolved.]

Platform and Infrastructure

- The output of the `show route forwarding-table family vpls multicast` command might display an unexpected output such as "rtinfo" with the multicast knob because this knob is supported only with inet and inet6 families and is not supported for te ISO, NTP, MPLS, UNIX, and VPLS families. The output of this command will be fixed in 10.1R1 to display the message: "Multicasting is not supported by the UNIX, ISO, NTP, MPLS, and VPLS protocols." [PR/235712: This issue has been resolved.]
- Reading the list of boot devices from the BIOS might fail once in hundreds or thousands of times due to an improper locking mechanism. [PR/461320: This issue has been resolved.]
- On T640 and TX Series routers with an outgoing interface on a GFPC, the interface might report LSIF errors or cell-mismatched errors after it receives an IPv6 packet with an invalid payload. The interface still accepts traffic, but discards all outgoing

packets. To recover, reboot the FPC on T640 and TX Series routers. However, if the IPv6 packets of the invalid payload are still transmitted, the problem will occur again. [PR/470219: This issue has been resolved.]

- When an aggregated SONET with a Cisco High-Level Data Link Control (HDLC) encapsulation is configured, a member link might not be marked as linkdown in the Packet Forwarding Engine if the remote end of the link is disabled. [PR/472677: This issue has been resolved.]
- The output of the `show arp` command does not show the entire demux interface identifier, making it difficult to determine with which specific demux subinterface a given ARP entry is associated. [PR/482008: This issue has been resolved.]
- The syslog usually logs data only when the per-fabric-stream counter increases. However, the syslog starts logging even if the counter value is not increasing. [PR/493384: This issue has been resolved.]
- The Source Class Usage (SCU) statistics counter value might drop occasionally when it is used with the accounting profile. [PR/493662: This issue has been resolved.]
- The traffic sent to ports on PB-4OC3-4OC12-SON-SFP PICs in an MX-FPC2 (sent above the configured bandwidth) might be dropped silently and non-deterministically. This uncontrolled traffic drop can lead to high priority traffic such as the PPP LCP being dropped. Depending on traffic conditions, this can cause a link configured for PPP to bounce indefinitely. [PR/493793: This issue has been resolved.]
- An issue occurs when one or more multicast routes (such as one or more `<S,G>s`) have received joins over an AE interface represented by two (or more) AE legs on separate Packet Forwarding Engines. In a Packet Forwarding Engine ASIC forwarding, the next hop shared by these multicast routes contains a list representing the two (or more) Packet Forwarding Engines. When this next hop list is no longer referenced by any active multicast route, it is not correctly freed and remains stranded in the Packet Forwarding Engine ASIC memory. This issue does not occur when the AE legs are all on the same Packet Forwarding Engine. [PR/494246: This issue has been resolved.]
- Due to excessive logging at the FPC, the E3 FPC Type 3 core dumps multiple times. [PR/494534: This issue has been resolved.]
- In certain cases, a configuration change can cause the backup Routing Engine to reboot. [PR/497290: This issue has been resolved.]
- On T Series routers with ES-FPCs, removing or adding flow-tap filters might trigger an FPC reboot. However, the other FPC types in the same system are not affected. [PR/499233: This issue has been resolved.]
- When a next-hop chain has multiple types of next-hop dependencies, including indirect next-hop, aggregate next-hop, and multiple unicast next-hops, during an aggregate link flap (down/up), a certain sequence of events from the kernel is expected by the Packet Forwarding Engine for the next-hop change and delete updates. However, during a quick link flap (down/up), in an extreme corner case, the Packet Forwarding Engine does not receive the expected sequence, and the FPC will crash. [PR/499315: This issue has been resolved.]
- On IQ2 PICs, when copy-plp is enabled under class of service, the DCU provides the wrong statistics. [PR/499378: This issue has been resolved.]

- The L2RW does not report an error when the required L2_pgm length is longer than what the hardware can support. [PR/501318: This issue has been resolved.]
- On an iChip platform, when the downstream multicast member link flaps, the Packet Forwarding Engine rarely has a chance to fail multicast next-hop handling. This can cause multicast traffic drops. [PR/501852: This issue has been resolved.]
- On a TX Matrix Plus router, if one of the two external RJ-45 links between a TXP-CIP and an LCC Control Board is broken, the router does not generate an alarm. [PR/508219: This issue has been resolved.]
- On some M, MX, and T Series routers, when a firewall filter is applied on the egress of an aggregate interface, packet loss might occur after adding, removing, or changing the service configuration on the egress side of the aggregate interface. As a workaround, deactivate and activate the output firewall filter on the aggregate interface. [PR/517992: This issue has been resolved.]
- When a socket connection between the Routing Engine and the FPC is reestablished, the FPC might run into a software crash because of an invalid counter being referenced. There is no workaround. [PR/525357: This issue has been resolved.]

Routing Protocols

- Deleting a logical system causes the routing protocol process to be stuck in an infinite loop. [PR/439000: This issue has been resolved.]
- The routing protocol process dumps core due to a soft assertion failed: "rt_notbest_sanity: Path selection failure" in rt_table.c. As a workaround, use the **bgp path-selection external-router-id** statement or the **bgp path-selection always-compare-med** statement. [PR/451021: This issue has been resolved.]
- When nonstop active routing (NSR) is running and BGP groups are added (eg a VRF with a BGP in it), the routing protocol process might crash. As a workaround, configure the new BGP groups after disabling the NSR. Reenable the NSR. [PR/487305: This issue has been resolved.]
- If there are enough routing instances with PIM configured, and there is enough IGMP/MLD join state present and a configuration change is made, a routing protocol process scheduler slip might occur. [PR/493062: This issue has been resolved.]
- On an unnumbered Ethernet interface in P2P mode, OSPF does not skip validation of the network mask received in the hello packets. This could result in a failure to bring up an adjacency on such interfaces while interoperating with other vendors. As a workaround, convert the interface to a regular numbered interface on both sides. [PR/493206: This issue has been resolved.]
- When l3vpn-composite-next-hop is configured, it should only be used by L3VPN routes. However, non-L3VPN routes are also able to use it. [PR/496028: This issue has been resolved.]
- In an NSR configuration, the backup Routing Engine can lose the connection to the active Routing Engine during a configuration commit. The problem occurs more often when the configuration includes a large number of routing instances. This is caused by the routing protocol process on the backup Routing Engine leaking file descriptors during commit synchronization. To recover, restart the

routing protocol process on the backup Routing Engine. [PR/506883: This issue has been resolved.]

- When the routing-instances *routing-instances-name* routing-options multipath vpn-unequal-cost equal-external-internal statement is configured, some VPN routes learned from different route reflectors can be shown as multipath. [PR/507236: This issue has been resolved.]
- Nonstop routing (NSR) does not work correctly if an automatic route distinguisher is used with a L2VPN routing-instance. [PR/513949: This issue has been resolved.]

Services Applications

- If the Juniper-Firewall-Attribute attribute in a RADIUS server configuration file names a policer that sets a bandwidth limit for Layer 2 Tunneling Protocol (L2TP) sessions but not an exclude-bandwidth limit, the bandwidth limit might not be set correctly. [PR/254503: This issue has been resolved.]
- On M Series routers (M120 and M320) with many service-sets configured with idp policies, kernel messages are seen in the messages file once traffic passes through these service-sets. These messages stop when the traffic is stopped. [PR/462580: This issue has been resolved.]
- A static route pointing to a destination is incorrectly added for a source NAT when a next-hop type service set is used. [PR/476165: This issue has been resolved.]
- MSDPC might crash while running a combination of SIP and other ALGs due to a possible double freeing of memory. [PR/491218: This issue has been resolved.]
- The SIP ALG on the services PIC might cause NAT port leaks in some call scenarios. [PR/491220: This issue has been resolved.]
- The `show services nat pool name` CLI filter does not have any effect. [PR/493820: This issue has been resolved.]
- Under certain conditions, the replication socket between two Routing Engines for the local policy decision function process (LPDFD) does not close properly. This results in high CPU consumption by the LPDFD. As a workaround, restart the local policy decision function process (LPDFD) on the master Routing Engine's restart local-policy-decision-function. [PR/495363: This issue has been resolved.]
- The l2tp on an M7i LNS crashes following an upgrade from JUNOS Release 9.3R1 to 9.6R2. [PR/498423: This issue has been resolved.]
- When using a NAT DCE RPC ALG on a services PIC, the PIC might crash while processing the binding request. [PR/510997: This issue has been resolved.]

User Interface and Configuration

- When an event policy is configured for an event with the attributes-match clause and if the event occurs without the attribute mentioned in the attributes-match clause, then the policy action gets executed. This behavior is wrong as the policy action should not be executed. [PR/421808: This issue has been resolved.]
- The wildcard apply groups do not work properly in JUNOS Release 9.1 and above. [PR/425355: This issue has been resolved.]

- The `deactivate` configuration statement is not blocked through the `deny-configuration` statement. [PR/488352: This issue has been resolved.]
- When commit scripts are used and the configuration contains a policy which uses an `apply-group` with a then action of “then community + EXPORT,” the commit fails. [PR/501876: This issue has been resolved.]
- On M10i, M120, M320, and MX Series routers with dual Routing Engines running JUNOS Release 9.4 or later, the `dfwd` process running on the backup Routing Engine might access the `/var/pdb/rdm.taf` file every 30 seconds, causing excessive writes to the hard disk drive. This problem does not occur when GRES is enabled. [PR/506691: This issue has been resolved.]

VPNs

- When different prefixes are advertised to the same source by different PE routers, an egress PE router is prevented from picking the lower prefix route for RPF when the PE advertising the higher prefix loses its route to the source. [PR/493835: This issue has been resolved.]
- While upgrading JUNOS Software with `l2circuit` configuration in the logical systems, the validation might fail with an “interface version mismatch” error. You can ignore this error and upgrade the JUNOS Software using the `no-validate` option. [PR/497190: This issue has been resolved.]
- When multipath is enabled in a routing instance with NG MVPN, the traffic might get dropped on the receiver PE. [PR/508090: This issue has been resolved.]

Previous Releases

Release 9.6R3

The following issues have been resolved since JUNOS Release 9.6R3. The identifier following the description is the tracking number in our bug database.

- Class of Service
- Forwarding and Sampling
- Interfaces and Chassis
- Layer 2 Ethernet Services
- MPLS Applications
- Network Management
- Platform and Infrastructure
- Routing Protocols
- Services Applications
- User Interface and Configuration
- VPNs

Class of Service

- After a graceful Routing Engine switchover, the packets might not be properly classified. [PR/452169: This issue has been resolved.]
- The BA classifiers for Ethernet VPLS over ATM traffic is used in the wrong queue. [PR/468936: This issue has been resolved.]
- On MX Series routers, the 64- to 67-byte L2 length packet might be dropped due to a tail drop on the 1-Gigabit Ethernet interface. This issue occurs because the performance restriction value was not increased after the intercomponent data rate was increased. [PR/469135: This issue has been resolved.]

Forwarding and Sampling

- The output firewall filter counter does not work when the firewall is configured for discard next hop. [PR/404645: This issue has been resolved.]
- When prefix actions are defined, after a graceful Routing Engine switchover, ISSU, or firewall restart, the error message "DFWD_CONFIG_WRITE_FAILED" might appear in the syslog. [PR/458573: This issue has been resolved.]
- When a configuration with a large number of terms is committed in a firewall filter, the error message: "DFWD_FW_COMPILER_EXIT_SIGNAL: dfwc exited with signal 11" displays and cores. [PR/465973: This issue has been resolved.]
- The `show firewall filter` CLI commands take a long time to display outputs with a large number of attachments. [PR/470794: This issue has been resolved.]
- Some ranges of burst sizes might result in unexpected packet drops when the traffic rates are close to the policing rate. Increase the burst size to resolve this problem. [PR/478659: This issue has been resolved.]

Interfaces and Chassis

- Under certain conditions, after a graceful Routing Engine switchover (GRES), the new master Routing Engine sends an invalid LACP frame. As a result, the aggregated interface fails. [PR/314855: This issue has been resolved.]
- The backup Routing Engine might fail to obtain mastership under the following cases:
 - Re0 gets stuck and doesn't reboot.
 - Due to some hardware problem, re0 loses its connectivity towards both the Control Board and the Packet Forwarding Engine.
 [PR/405412: This issue has been resolved.]
- The bandwidth on any logical interface configured on a physical interface shows a value that is greater or smaller than the speed of the respective physical interface. [PR/426469: This issue has been resolved.]
- With the `show interfaces extensive` command, some interfaces might not display the correct value for the Oversized Frames counter. [PR/437176: This issue has been resolved.]

- On an MX960 router, when more than eight Dense Port Concentrators (DPCs) (including unconfigured DPCs) are loaded, the output of the **show interface extensive** command can be very slow if the source class usage destination class usage (SCU/DCU) is configured for some units. [PR/449034: This issue has been resolved.]
- Interrupts occurring from links (non-zero) that are not configured or enabled in the PIC due to a hardware issue in the DFPGA cause the syslog to overload and eventually lead the FPC to core. [PR/455877: This issue has been resolved.]
- When a MAC address is learned, it should be added in both the sa-mac and da-mac tables. The log message "remove mac entry failed for serv id" is continuously logged for IQ2 PICs because the MAC address that isn't added to one table is not deleted from the other table. [PR/459890: This issue has been resolved.]
- The master Routing Engine fails to establish a connection with the backup Routing Engine due to an autonegotiation issue with the em1 interface. [PR/461469: This issue has been resolved.]
- For Tri-Rate Enhanced Dense Port Concentrator (DPC) interfaces, the link LED does not reflect the correct status when the interface speed is set to auto. [PR/466588: This issue has been resolved.]
- The error "arp_update_iff_vrrp: IFF ae11 doesn't have a vrrp group configured" occurs when the native-vlan-id is configured with the Virtual Router Redundancy Protocol (VRRP). [PR/468167: This issue has been resolved.]
- The JUNOS Software fails to get EEPROM data when the **show chassis hardware** command is used. [PR/468459: This issue has been resolved.]
- When an untagged aggregated Ethernet interface is configured with LACP and GE IQ2 PICs as the child interface, the input packet count might be constantly decremented to zero when no data packets arrive on the interface. The decrease in packet count is equal to the incoming LACP packet count. [PR/471177: This issue has been resolved.]
- Commits will fail for ATM interfaces when the virtual path identifier (VPI) is configured in atm-options promiscuous-mode. As a workaround, avoid using the configurations that DCD erroneously rejects. [PR/471905: This issue has been resolved.]
- Enabling and disabling LMI keepalives on a Frame Relay encapsulated interface while the physical interface is down, might result in the logical interfaces remaining down when the physical interface is restored. Deactivate and activate the logical interfaces to recover connectivity. [PR/472688: This issue has been resolved.]
- With a default configuration, when a Tri-Rate copper small form-factor pluggable (SFP) transceiver installed in a DPCE-R-20GE-2XGE board is replaced with a SFP-LX/SFP-SX, the link stays down. Activate and deactivate the SFP to restore the link. [PR/473127: This issue has been resolved.]
- On an M320 router, the 4x STM-1 1x STM-4 SFP PIC (PB-4OC3-1OC12-SON-SFP) currently supports only two ports (0 and 2) when configured for eight queues per port on an E3 FPC. [PR/475008: This issue has been resolved.]

- Multiple vmcore/kernel and ksyncd core files might be generated when graceful Routing Engine switchover (GRES) is enabled and the router is being used for subscriber management. [PR/480734: This issue has been resolved.]

- When a DPC restarts, a large number of routes (about 700,000 simple IPv4 routes) in the forwarding table are still learned through another DPC. The sync process between the Routing Engine and the Packet Forwarding Engine takes too long, and the Routing Engine will restart the FPC. This repeats endlessly.

To restore the service and stop the DPC boot loop, restart the chassis process or the routing process. [PR/481164: This issue has been resolved.]

- On a 4x CHOC3 SONET CE SFP PIC and 12x T1/E1 CE PIC, if a T1 or E1 interface is deleted and re-created, the T1 or E1 interface that is connected to a 4x CHOC3 SONET CE SFP PIC or 12x T1/E1 CE PIC might observe a framing error and traffic might not pass.

As a workaround, deactivate the E1 interface encapsulation, then activate the E1 interface encapsulation after the T1 or E1 interface is deleted and re-created on a 4x CHOC3 SONET CE SFP PIC or 12x T1/E1 CE PIC. This will make the framing error disappear. [PR/482491: This issue has been resolved.]

- Under certain conditions, when aggregate interfaces are used, and the member links are located on more than one FPC, multicast traffic will not use one or more of the aggregate child links. This can happen after an FPC reboot.

If the aggregate member links are located on the same FPC, this problem is not triggered. To recover from this condition, deactivate and activate the aggregate interface. [PR/484007: This issue has been resolved.]

- The logical unit of a Gigabit Ethernet interface might show less than 1000 Mbps of bandwidth even if there is no speed configuration under the physical interface. As a workaround, manually set the bandwidth on the logical interface. [PR/485840: This issue has been resolved.]

- On an M20 router with an LS PIC, the backup Routing Engine kernel might core at `rnh_index_alloc`. [PR/486646: This issue has been resolved.]

- Traffic might be sent out on a child link of an Aggregated Ethernet (AE) bundle even when it is not in the Collecting-Distributing Link Aggregation Control Protocol (LACP) state if and only if the following conditions are met:

- The remote end configured one link to be primary and another to be backup.
- On the System Under Test (SUT), a unit of the AE bundle is disabled, then enabled.

As a workaround, deactivate and activate the child link that is not in the Collecting-Distributing LACP state. [PR/487786: This issue has been resolved.]

- On an IEEE 802.1ag CFM, when the loss-threshold is configured to 256, it displays a "0." [PR/491422: This issue has been resolved.]
- An outer virtual LAN tag is not added in a provider edge-customer edge link when VPLS traffic arrives with an MPLS value of 2, 3, 4, or 5. However, VPLS traffic with a value of 0, 1, 6, or 7 does not have this issue. [PR/495555: This issue has been resolved.]

Layer 2 Ethernet Services

- On MX960 platforms, the up and down fan speed transitions occur at the same temperature, which might cause excessive transitions and messages to be logged. [PR/462044: This issue has been resolved.]
- In a combo DPC, the physical link stays up when an interface with the SFP-T is disabled. However, port 0 of the combo DPC is not impacted by this issue. [PR/477848: This issue has been resolved.]
- On an MX Series router, the DHCP ACK messages are dropped when a client Rebind request is processed by a different DHCP server. This issue might occur in an environment where the provider has multiple DHCP servers for redundancy purposes. [PR/487138: This issue has been resolved.]
- The family ISO MTU configured explicitly under the IRB interface logical unit will decrement by three if you change the interface MTU on the interface that belongs to the same bridge domain. [PR/493209: This issue has been resolved.]

MPLS Applications

- When a point-to-point LSP has a primary path and some secondary paths that are experiencing signaling problems (when CSPF computation is fine, but signaling keeps failing and retrying) except for one secondary path, and a make-before-break signaling is performed for that secondary path (that is up), due to auto-bandwidth or path re-optimization, it is possible that the LSP's RSVP source port (LSP ID) space might wrap. When this happens, the source port of that secondary path will be allocated to another path, causing the PSB of the secondary path to be associated with the other path, eventually leading to a routing protocol process (RPD) crash.

Similarly, the branch ID and the source port ID (LSP ID) of a point-to-multipoint LSP might wrap in some make-before-break and signaling retry situations, causing the same RPD crash. [PR/265242: This issue has been resolved.]
- Configuration of a non-existent IP address in MPLS for a label-switched path could result in a memory leakage in the routing protocol process. [PR/459254: This issue has been resolved.]
- If both OSPF and IS-IS update the traffic engineering database (TED) on the same traffic engineering link, it might take some time for OSPF to update the traffic engineering database with the new MPLS administrative group (affinity) after the administrative group configuration is changed. [PR/465953: This issue has been resolved.]
- When replying to an MPLS RVSP traceroute with an MPLS Echo Reply packet, the label-switching router (LSR) populates the downstream IP address value with 0.0.0.0 instead of the actual IP address. [PR/466049: This issue has been resolved.]
- The JUNOS Software does not recognize the LDP TLVs with Ignore (U) or Forward (F) bits set as it is not in conservative retention mode. [PR/467164: This issue has been resolved.]

- When a large number (more than 100) of NGEN-MVPN P2MP LSPs based on an LSP template are active, the routing protocol process might crash if the LSP template is deleted and added back. [PR/477376: This issue has been resolved.]
- Under some circumstances where LDP is enabled, a memory leak might occur where the routing protocol process does not free up memory. [PR/493885: This issue has been resolved.]

Network Management

- The `snmpwalk` on `ipNetToMediaPhysAddress` might show some ARP entries missing from the output when displayed using the `show arp` command. [PR/453855: This issue has been resolved.]
- A problem with the IPv6 n2m add routine causes the `mib2d` to fail at the `vlogging_event`. [PR/472453: This issue has been resolved.]
- The `snmpwalk` on `jnxFWCounterDisplayName` might miss certain policer counters of firewall filters applied with respect to logical interfaces (subinterfaces). [PR/485477: This issue has been resolved.]

Platform and Infrastructure

- On M320 and T Series routers, a process 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: This issue has been resolved.]
- Under some circumstances, the interface process (physical interface) might interfere with the operation of an LSI interface. [PR/102431: This issue has been resolved.]
- When certain FPCs (T1600-FPC4-ES, T640-FPC4-1P-ES, T640-FPC1-ES, T640-FPC2-ES, and T640-FPC3-ES) receive corrupted cells via 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.]
- A large amount of syslog messages generated by the Packet Forwarding Engine components might take more CPU cycles of the S-boards to process the syslog queue. This causes them to restart. [PR/454070: This issue has been resolved.]
- The Packet Forwarding Engine will change the dependent next hop to discard when an IFF is deleted. [PR/459781: This issue has been resolved.]
- The current CLI knob in the `[edit system saved-core-files]` hierarchy level that sets the `sysctl` variable `debug.ncores` does not set values greater than 10 despite having a value range from 1 to 64. [PR/466461: This issue has been resolved.]
- Using link flaps with many VRF prefixes along with the statement `l3vpn-composite-nexthop` might result in jtree corruption which triggers traffic black-holing. [PR/468584: This issue has been resolved.]
- After upgrading from JUNOS Software Release 9.3 to Release 9.5, the timestamps in the log files show the UTC time instead of the local time corresponding to the specified time zone. [PR/469175: This issue has been resolved.]

- On MX Series routers with FPC under flow-control condition, the keepalive packets are dropped from the ASIC without a trace. [PR/470334: This issue has been resolved.]
- An FPC might stop forwarding traffic when an aggregate interface flaps and the router uses per-prefix load balancing (default configuration) for some prefixes. A more likely scenario under which this issue can occur is when an aggregate interface is configured with just a single link (that flaps), and per-prefix load balancing is used.
As a workaround, use a load balancing per-packet policy for all prefixes (per-flow load balancing) and/or do not have aggregate interfaces flap. [PR/477326: This issue has been resolved.]
- With JUNOS Release 9.3 or later, configuring policer or SCU/DCU on interfaces belonging to FPC-ES might cause memory corruption which leads to either traffic loss, or the FPC to restart unexpectedly. [PR/481185: This issue has been resolved.]
- If a duplicate IPv6 address is configured, every ICMP6 packet received (icmp request, icmp neighbor solicitation, or icmp neighbor advertisement) will trigger an mbuf leak. Such a duplicate address configuration might not be noticed at the VRRP backup router which is not used for data forwarding. Correcting the configuration, and deactivating and activating the interface will stop the mbuf leak. [PR/482202: This issue has been resolved.]
- On T Series routers with ES-FPCs, removing or adding flow tap filters might trigger an FPC reboot. However, the other FPC types in the same system are not affected. [PR/499233: This issue has been resolved.]

Routing Protocols

- When an IGMP join occurs for a group from snooped and non-snooped BDs, or both the aggregated interface and irb interface are multicast downstreams, the traffic might get black-holed. [PR/441639: This issue has been resolved.]
- An interface with a higher priority is not elected as a PIM DR when a default priority is not used on one side. [PR/453561: This issue has been resolved.]
- For a certain error condition during negotiation with a very old router, the sending of the 4-byte AS capability is not consistent with the sending of the other capabilities. [PR/462930: This issue has been resolved.]
- Whenever a mastership switchover is done, the PPMD might pull the Packet Forwarding Engine connections before transitioning into the slave, and before receiving the SIGHUP, which could result in the BFD sessions being incorrectly marked as down. [PR/465911: This issue has been resolved.]
- A core might occur if the "rib-group" configuration references the inetflow table of a routing instance, `ri-name.inetflow.0`, and the routing instance does not exist in the configuration. [PR/467332: This issue has been resolved.]
- If a router modifies the next-hop protocol to self (for example, using an export policy with next-hop-self) on a peer group containing "internal" peers, and nonstop routing is configured on the router, the routing protocol process might send duplicate updates to the peers in this peer group during a Routing Engine switchover. [PR/468505: This issue has been resolved.]

- If a reject route is present for the address of a Multicast Source Discovery Protocol (MSDP) SA originator, the routing protocol process will crash. [PR/469142: This issue has been resolved.]
- Due to the routing protocol process startup time discrepancy between the master and the backup Routing Engine, overload state might occur after the Routing Engine switchover when the overload timer has expired on the previous master Routing Engine. When ISSU is configured, the new master Routing Engine starts up and enters into the overload mode. [PR/472408: This issue has been resolved.]
- When nonstop routing is configured on a router, the routing protocol process might restart with a core dump. [PR/472701: This issue has been resolved.]
- Rapid changes in the status of a particular route while the system is under high load could cause the routing protocol process to restart unexpectedly. [PR/473517: This issue has been resolved.]
- When the routing protocol process (rpd) fails, after a routing protocol process restart, the process might be unable to install new LSI logical interfaces. The following error is returned: "ENOMEM." [PR/473774: This issue has been resolved.]
- If the routing options forwarding-table indirect-next-hop knob is set in the configuration, every commit that changes any parameter related to routing will result in several minutes of 90 percent or more of CPU load by the routing protocol process. [PR/475117: This issue has been resolved.]
- During an ISSU upgrade, the BGP session might flap due to differences in the negotiation of keepalive messages between versions. [PR/476285: This issue has been resolved.]
- After a mastership switchover, incorrect BFD packets might be sent out due to stale information within the ppmmd. This might result in the BFD sessions flapping repeatedly. [PR/478447: This issue has been resolved.]
- Under certain circumstances, the Juniper Networks PIM implementation might send (S,G,rpt) prune messages towards the RP too early after receiving the (S,G,rpt) prune message from a downstream router. [PR/478589: This issue has been resolved.]
- The routing protocol process (RPD) CPU usage might be high if both BGP multipath and family inet-mpvn are configured under BGP. [PR/479574: This issue has been resolved.]
- When running PIM and a link flap occurs, the routing protocol process might core. [PR/480422: This issue has been resolved.]
- The MVPN c-multicast traffic is duplicated onto the LAN segment since the interface mismatch is not processed within the PIM. Interface mismatch is needed to trigger an assert to prevent traffic duplication. As a workaround, configure PIM under the main instance. [PR/481467: This issue has been resolved.]
- Whenever a graceful Routing Engine switchover (GRES) is performed, the BMP header for the consequent updates might get corrupted until the BMP session is deactivated and activated. [PR/486068: This issue has been resolved.]
- The output of the **show igmp interfaces** command might display the configured IGMP query-interval value incorrectly in the output. [PR/488146: This issue has been resolved.]

- The routing protocol process might core frequently because of malformed BGP updates generated by the JUNOS Software. This could be because of the total length and the path attribute length. [PR/489891: This issue has been resolved.]
- The MPLS LSPs are not advertised as links into the non-backbone OSPF areas, even though they are configured to be advertised. [PR/491692: This issue has been resolved.]
- The PIM running in the main instance might stop working if the PIM is configured in a no-forwarding routing instance. [PR/492017: This issue has been resolved.]
- When running on an unnumbered Ethernet interface in P2P mode, the OSPF does not skip validation of the network mask received in hello packets. This could result in a failure to bring up an adjacency on such interfaces when interoperating with other vendors. As a workaround, convert the interface to a regular numbered interface on both sides. [PR/493206: This issue has been resolved.]
- With NSR configured, the routing protocol process might crash in situations where the IPv6 PIM neighbor link-local address is changed as a consequence of a neighbor interface MAC address change. [PR/493783: This issue has been resolved.]

Services Applications

- During packet transmission of the L2TP/MLPP bundle with no fragmentation, the service loses track of the sequence number and starts dropping most of the input L2TP packets as "fragment out of range." [PR/430296: This issue has been resolved.]
- When two different filters with different source-port values are configured in the `X-JTap-Cdest-Source-Port` parameter of the filter specification, an "Invalid filter specification" error occurs. [PR/447855: This issue has been resolved.]
- MS-PIC might crash while handling Real-Time Streaming Protocol (RTSP) flows. [PR/455649: This issue has been resolved.]
- On MX960 routers, the NAT "ports in use" count displayed using the `show services nat pool detail` command is greater than the SFW flow count displayed using the `show services stateful-firewall flows count` command. [PR/466506: This issue has been resolved.]
- When a SIP malformed packet that is not compliant with RFC 2543 in ch.6.40 is received by the SIP Alg, the service PIC might restart. [PR/467600: This issue has been resolved.]
- The service DPCs might crash during conversation timeout cleanup for the DCE-RPC. [PR/475436: This issue has been resolved.]
- When a malformed RTSP packet not conforming to RTSP RFC syntax is processed by the RTSP Application Layer Gateway (ALG) within the service PIC (or Service DPC), the PIC might dump a core. [PR/476321: This issue has been resolved.]
- Via header translation might be incorrectly performed by the SIP ALG when it contains only an IP address and no port. [PR/482998: This issue has been resolved.]

- The SIP ALG does not translate the route header properly, which leads to the SIP calls being dropped after 20 seconds. [PR/483014: This issue has been resolved.]
- The SIP parser might drop 200 “OK for REGISTER” messages if the contact has multiple entries. [PR/483030: This issue has been resolved.]
- When SIP ALG is enabled on ASPIC, MSPIC, or MSDPC, the PIC could crash while freeing the Via header NAT port. [PR/490329: This issue has been resolved.]
- Under certain conditions, the replication socket between two Routing Engines for the local policy decision function process (LPDFD) does not close properly. This results in high CPU consumption by the LPDFD. As a workaround, restart the local policy decision function process (LPDFD) on the master Routing Engine’s ‘restart local-policy-decision-function’. [PR/495363: This issue has been resolved.]

User Interface and Configuration

- When the syslog configuration for forwarding messages to a remote host has a source address configured, the messages are not filtered by regular expressions. [PR/446140: This issue has been resolved.]
- When jcs:syslog() is used in an event script, messages do not appear until another system application sends a syslog message. [PR/449778: This issue has been resolved.]
- After an NSM server restarts, the Mgd process is at a high CPU utilization of 95 percent or more. [PR/455166: This issue has been resolved.]
- The router does not return the username in the accounting packet sent to the RADIUS server. The following issues have been noticed:
 - The acc-start uses the “remote” username despite the real username being available.
 - The interim-update has no username.
 - The stop message has no username.

[PR/472704: This issue has been resolved.]

- If a configuration contains both a system backup-router destination and a routing-options static route (that did not include the “retain” flag), and they use the same prefix, a warning message is generated for every commit: “warning: n.n.n.n/m is also configured as backup-router destination, turn on retain flag to ensure proper functionality.”

This warning causes the routing configuration to be completely re-evaluated for each commit. This causes slow commits, and high CPU load by the routing protocol process during and after the commit. As a workaround, add the “retain” flag to the indicated static route. [PR/473204: This issue has been resolved.]

VPNs

- Configuring a forwarding-cache threshold under a routing instance for NG-MVPN might not produce the expected behavior and might not limit the number of forwarding cache entries. [PR/438164: This issue has been resolved.]
- The routing protocol process might crash if the Routing Engine containing the auto-RD configuration for L2VPN or VPLS routing instances is rebooted. [PR/469847: This issue has been resolved.]
- On an MX960 router, the VPLS instance might not learn the remote CE MAC address when the `clear vpls mac-address` command is used. [PR/476020: This issue has been resolved.]
- P2MP LSP cannot be recovered when the P router (which is also configured as a BGP reflector) goes down. [PR/481441: This issue has been resolved.]

Release 9.6R2

The following issues have been resolved since JUNOS Release 9.6R2. The identifier following the description is the tracking number in our bug database.

- Class of Service
- Forwarding and Sampling
- General Routing
- High Availability
- Interfaces and Chassis
- Layer 2 Ethernet Services
- MPLS Applications
- Network Management
- Platform and Infrastructure
- Routing Protocols
- Services Applications
- User Interface and Configuration
- VPNs

Class of Service

- After the aggregate chassis configuration is deactivated then activated, the classifier might not be properly applied on the aggregate interfaces. Deactivate and activate class of service to fix the problem. [PR/442240: This issue has been resolved.]
- The "Dropped packets" counter continues to increment on the ingress queue with the respective egress queue "Dropped packets" counter, even if there is only egress traffic on an IQ2 GE interface. [PR/453253: This issue has been resolved.]

- With JUNOS Release 9.3R4.4 or 9.4R3.5 in an MPLS environment, when T Series or TX Series routers perform PHP, unpredictable values might be written to the IP TTL field of the egress IP packets. This occurs only on T Series Enhanced Scaling FPCs (FPC Type 3-ES, FPC Type 4-ES) and when the router performs PHP. [PR/463989: This issue has been resolved.]
- M120 and MX Series routers might have incorrect transmit-rate values programmed for interface output queues. [PR/467103: This issue has been resolved.]

Forwarding and Sampling

- When JUNOS Software is upgraded from Release 8.x to Release 9.1 and above, high CPU utilization and memory allocation failures might occur when running sampled process. This issue occurs because the sampled process takes a larger memory due to the size of the AS-path record changed from 16 bit to 32 bit. [PR/448521: This issue has been resolved.]
- When one **fast-update-filter** term contains a range match on DSCP and another term has a DSCP match for a single value within that range, the addition of the second term will fail with a “term conflict error”. The DSCP range can either be an explicit range or the implicit default range (this is the case where DSCP is in the match-order but not in the term's “from” stanza). As a workaround, configure multiple terms so that the DSCP ranges and values do not overlap. [PR/460632: This issue has been resolved.]

General Routing

- The **show helper statistics** command displays the error: “dhcpd subsystem is not running”, even when the bootp helper configuration is successfully configured and committed. [PR/445240: 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.]
- After an unified in-service software upgrade (ISSU) on the MX Series router, you might see a kernel database replication error, 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.]
- After an unified in-service software upgrade (ISSU) on an MX Series router, there might be fluctuations in output traffic due to flooding. The MAC tables on the DPCs are out-of-sync, as no updates are received with regard to the MAC addresses. This issue is present only in JUNOS Releases 9.5 and 9.6. [PR/461822: This issue has been resolved.]

Interfaces and Chassis

- The OC192 XFP might display the "XFP read fail, retry for 1 times" message at random intervals. This is a cosmetic issue and does not affect the functionality of the interface. [PR/262883: This issue has been resolved.]
- The MTU size cannot be changed on fxp0/fxp1 interfaces. [PR/419379: This issue has been resolved.]
- The bandwidth on any logical interface configured on a physical interface shows a value that is greater or smaller than that of the speed of the respective physical interface. [PR/426469: This issue has been resolved.]
- The JUNOS CLI allows invalid combinations of **atm-l2circuit-mode** encapsulation on atm-ce interfaces. The consequence of using the wrong encapsulation combination in the CLI is that the incorrect setting is ignored and the ATM pseudowire behaves as if the configuration does not exist. No error message will be displayed. As a workaround, configure the correct **atm-l2circuit-mode** encapsulation. [PR/437253: This issue has been resolved.]
- Performing an RLSQ switchover after multiple graceful Routing Engine switchovers might result in one of the Routing Engines displaying the db prompt. [PR/450570: This issue has been resolved.]
- Referencing nonexistent loopback interfaces when dynamically creating LANs or stacked VLANs on MX Series routers is not supported. Referencing a nonexistent loopback interface in a VLAN or stacked VLAN dynamic profile can result in the device control process (DCD) failure upon subscriber login. [PR/457710: This issue has been resolved.]
- With a specific SFP-T set to "disable" when installed in an DPCE 20x 1GE + 2x 10GE R or DPCE 20x 1GE + 2x 10GE R EQ, the router displays the following messages in the message log :

```
fpc4 ge-4/0/1: an_link_state_chg, error=20 :xeth_mac_program_sfp_phy
fpc4 PQ3_IIC(WR): no target ack on byte 0 (wait spins 3)
fpc4 PQ3_IIC(WR): I/O error (i2c_stat=0xa3, i2c_ctl[1]=0xb0, bus_addr=0x56)
fpc4 XETH(4/0): Failed to set SFP Module(0x32, 0x56)to start addr 0x0
fpc4 XETH(4/0): i2c mux controlling toxic SFP reset
fpc4 ge-4/0/2: phy_an_control, error=20 :xeth_an_chg_ctrl
fpc4 AN: an_chg_ctrl, ifd=ge-4/0/2, an_enable=1, err=20 :an_apply_cfg
fpc4 AN: an_apply_cfg, an_enable = 1 event = 5 ifd=ge-4/0/2, error=20 :
an_unspec_handle_event
```

Once the **disable** statement is removed from the configuration, the port remains down. As a workaround, remove and insert the specific SFP-T, or restart the DPC. [PR/458774: This issue has been resolved.]

- When a VRRP inherit is configured on two or more interfaces with one of the interfaces in link down state, the **show vrrp** command displays only those interfaces that are in the link down state. [PR/459630: This issue has been resolved.]
- Commits will fail for ATM interfaces when the virtual path identifier (VPI) is configured in atm-options promiscuous-mode. As a workaround, avoid using the configurations that DCD erroneously rejects. [PR/471905: This issue has been resolved.]

Layer 2 Ethernet Services

- When the DHCPv6 clients are bound, the libstats show an error message to the `jdhcpd` requests, even though the client binds successfully. This error message can be ignored as it does not affect the binding of the DHCP subscriber. [PR/435334: This issue has been resolved.]
- Applying the `clear dhcp server binding` command on a specific IPv4 address does not work when DHCPv6 is configured. [PR/459387: This issue has been resolved.]
- If member links of an AE are on different DPCs, the order of the `ppmd` stats reply might get changed in the distribution mode, which can cause the LACPD to core. [PR/460900: This issue has been resolved.]
- Multicast packets received on an AE interface that is part of an IRB are counted twice, once for the bridged packet and once for the routed packet. [PR/461923: This issue has been resolved.]

MPLS Applications

- When an uplink on a PE router is deactivated, the MPLS LSP BFD session over this link might not switch to other uplinks. [PR/454071: This issue has been resolved.]
- When an MPLS traceroute is executed in a downstream mapping TLV (TLV 2), the reply packet contains misleading values because of an MPLS OAMD error. [PR/454796: This issue has been resolved.]
- When hop-limit is enabled using the `protocols mpls label-switched-path` command along with FRR, the detour path might not be established successfully. [PR/462074: This issue has been resolved.]
- The JUNOS Software does not recognize the LDP TLVs with Ignore (U) or Forward (F) bits set as it is not in conservative retention mode. [PR/467164: This issue has been resolved.]

Network Management

- The `SnmpTrapEnterprise` variable is being added to generic traps such as link up, link down, and authentication failure. [PR/453092: This issue has been resolved.]

Platform and Infrastructure

- Following an FPC reset, the next-hop route pointing to the service PIC interface running RPM might be incorrect. [PR/438599: This issue has been resolved.]
- Traffic flow halts abruptly from an RSLQ interface when the interface is deactivated and activated. [PR/440564: This issue has been resolved.]
- Some catastrophic events like the FPC going down can corrupt the next-hop databases. In such cases, panic is caused by stack overflow or recursion, and can result in a core dump. [PR/448074: This issue has been resolved.]

- On TX Matrix Plus routers, when the LCC's Routing Engine reboots, the `tnp.sntpd` might stop working due to bogus NTP query packets from the rebooting LCC's Routing Engine. [PR/450217: This issue has been resolved.]
- MX Series tunnel interfaces configured on the DPC show traffic incorrectly on other interfaces. [PR/450844: 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. As a result, it reboots on just one single-bit-error event. [PR/457157: This issue has been resolved.]
- On M320 or MX Series routers with E3-FPC, every change or removal of the local interface address might trigger a jtree memory leak of 16 bytes. This leak occurs only when a loopback firewall filter is configured. The jtree memory is also not freed even when you remove the loopback filter. [PR/457717: This issue has been resolved.]
- When certain IP packets are received and transmitted via integrated routing and bridging (IRB) into a VPLS instance and these packets require fragmentation, the interface might stop transmitting. It requires a DPC reset to recover from this state. [PR/458423: This issue has been resolved.]
- On M Series and T Series routers, the Packet Forwarding Engine might core dump or assert without a core dump upon an Ethernet transmit ring buffer overflow condition. [PR/462934: This issue has been resolved.]
- Using link flaps with many VRF prefixes along with the `l3vpn-composite-nexthop` statement might result in jtree corruption which triggers traffic black-holing. [PR/468584: This issue has been resolved.]

Routing Protocols

- The routing protocol process might restart if PIM is configured to run on unnumbered interfaces. [PR/295319: This issue has been resolved.]
- With BGP NSR configured along with route-flap damping, when damping occurs for some prefixes at the same time as the initial state synchronization between the master and the backup rpd, the backup rpd could core after the initial state synchronization completes. [PR/312098: This issue has been resolved.]
- On routers running OSPF and advertising indication LSA for a DC-incapable neighbor, the routing might become corrupted. This will cause the RPD to crash when the LSA gets purged. [PR/406276: This issue has been resolved.]
- On a scaled BGP configuration with NSR configured, BGP might misinterpret socket fullness and close the session instead of retrying. [PR/443507: This issue has been resolved.]
- When a router has a flapping route and its next-hop is also changing, the routing protocol process might core. [PR/448629: This issue has been resolved.]
- When a switchover is triggered due to RPD thrashing, the BFD in the new master remains at admin-down state forever. [PR/451211: This issue has been resolved.]
- When nonstop active routing (NSR) is enabled, alternate paths for BGP prefixes with identical attributes might not be copied to the backup Routing Engine upon a Routing Engine switchover or in other situations when the backup Routing

Engine needs to learn routing updates from the master Routing Engine.
[PR/458402: This issue has been resolved.]

- When the BGP NSR is configured along with sampling (under forwarding-options sampling), duplicate updates for some prefixes could be sent during the Routing Engine switchover. [PR/458669: This issue has been resolved.]
- A failure might occur if the "rib-group" configuration references the inetflow table of a routing instance, `ri-name.inetflow.0`, and the routing-instance does not exist in the configuration. [PR/467332: This issue has been resolved.]
- If a reject route is present for the address of a Multicast Source Discovery Protocol (MSDP) SA originator, the routing protocol process will crash. [PR/469142: This issue has been resolved.]

Services Applications

- During packet transmission of the L2TP/MLPP bundle with no fragmentation, the service loses track of the sequence number and starts dropping most of the input L2TP packets as "fragment out of range." [PR/430296: This issue has been resolved.]
- While using the dynamic endpoint configuration for IPsec services, and dedicated interfaces on the same gateway address, only the first interface configured in a given routing instance is able to forward traffic. [PR/448498: This issue has been resolved.]
- When using Service PIC or Service DPC with stateful firewall and NAT service, some specific SIP traffic might cause the PIC to fail. [PR/459378: 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: This issue has been resolved.]
- When there are a lot of flows, the CLI session hangs when the `show services local-policy-decision-function flows interface interface` command is issued. [PR/447774: This issue has been resolved.]
- If a commit-script uses the get-configuration API element to request a copy of the configuration, the JUNOS Software will hang for hours during bootup. When it finally starts, there are no committed configurations. It is recommended to use the commit scripts to check the version of the JUNOS Software before using this API element. Do not perform this API request on an unfixed version. [PR/452398: This issue has been resolved.]
- There is a memory leak in the eventd process upon executing event scripts. This leak is fairly slow (8000 to 24,000 per script execution). [PR/457989: This issue has been resolved.]

VPNs

- After the ingress PE router for a next-generation MVPN instance performs a GRES event, the egress PE routers could fail to install a new forwarding state for the multicast traffic. As a workaround, clear the BGP session on the ingress router to restore traffic to all egress routers. [PR/441392: 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 process (RPD) might restart, creating P2MP LSP paths. [PR/442544: This issue has been resolved.]

Release 9.6R1

The following issues have been resolved since JUNOS Release 9.5R3. The identifier following the description is the tracking number in our bug database.

- Class of Service
- Forwarding and Sampling
- High Availability
- Interfaces and Chassis
- Layer 2 Ethernet Services
- MPLS Applications
- Network Management
- Platform and Infrastructure
- Routing Protocols
- Services Applications
- Subscriber Access Management
- VPNs

Class of Service

- In the cosd logs for JUNOS Release 9.4R1, "entries" is misspelled as "enteries." [PR/439993: This issue has been resolved.]
- When an Intelligent Queuing PIC is taken offline and 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 the PIC comes back online. [PR/444543: This issue has been resolved.]
- When a classifier is applied on a services PIC logical interface, a commit warning is issued stating that the classifier is not supported on this interface. [PR/448913: This issue has been resolved.]

Forwarding and Sampling

- On M320 and T Series routers, 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: This issue has been resolved.]
- On T Series routers, if an ingress firewall is configured to drop all incoming multicast packets. The discarded multicast packets are sent to the Routing Engine incorrectly. This causes a high utilization of the CPU (50 percent) on the FPC. [PR/239268: This issue has been resolved.]
- When configuring routing instances in a firewall filter, the router will display a warning message “Warning: statement ignored: unsupported platform.” [PR/421765: This issue has been resolved.]
- Upon changing policers on an Aggregated Ethernet interface, the DPC might reboot. [PR/431635: 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.]

Interfaces and Chassis

- On the Channelized STM-1 with QPP PIC, error monitoring for CRC and Frame Errors might not work as expected. [PR/39440: This issue has been resolved.]
- 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: This issue has been resolved.]
- 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: This issue has been resolved.]
- The XML output is not correct when the VRRP track interface is configured. [PR/414734: This issue has been resolved.]
- Under some conditions, if an interface flaps for an interval less than the hold-down time value configured, an interface might stop forwarding even though it shows as being UP. As a workaround, enable traffic monitoring on the interface or enable and disable the interface. [PR/423065: This issue has been resolved.]
- Upon changing policers on a Aggregate Ethernet interface, the DPC might reboot. [PR/431635: 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.]
- Too many ATM2 error interrupts might cause the FPC to panic. [PR/438073: This issue has been resolved.]
- When you configure the `payload port-data` statement at the `[edit family mpls hash-key]` hierarchy level on M120, MX Series, or M320 routers with E3 FPCs, the hashing algorithm might not take the port-data values into account. [PR/442223: This issue has been resolved.]
- On M Series routers, BGP sessions flap when any configuration change (even an irrelevant one) happens. As a workaround, make the difference between the configured MRRU and MTU to be greater than eight. [PR/442688: This issue has been resolved.]
- If VRRP tracks a cloned route, then the cloned route will always be treated as down. This is because the unicast cloned routes do not get added to the routing table. [PR/446408: This issue has been resolved.]

Layer 2 Ethernet Services

- When you configure graceful Routing Engine switchover (GRES) on MX Series routers, the Switch Interface Board (SIB) might not initialize if you reboot both Routing Engines simultaneously or reboot a router with only one Routing Engine installed. [PR/408359: This issue has been resolved.]

MPLS Applications

- 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: This issue has been resolved.]
- 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: This issue has been resolved.]
- The `mplsResourceTunnelTable` reports bandwidth in bps instead of kbps. [PR/432716: This issue has been resolved.]
- MPLS LSP auto-bandwidth adjustment might stop working while RSVP signals for the path; either optimization is initiated or the LSP goes down. [PR/438157: This issue has been resolved.]

Network Management

- When the SNMP get response is larger than 9 KB, a "Message too long" log is reported but no SNMP gets a response failure with a code "tooBig" sent back to the source. [PR/389559: This issue has been resolved.]
- tcpdump might report a max-response-time within IGMP in seconds while it is presenting units of 1/10th of a second. [PR/424618: This issue has been resolved.]

Platform and Infrastructure

- On T Series routers, the commit operation succeeds when you include the `no-labels` statement at the `[edit forwarding-options hash-key family mpls]` hierarchy level, but MPLS labels are still included in the hash key. [PR/80334: This issue has been resolved.]
- After an ISSU software upgrade on the MX Series router, you might see a kernel database replication error, 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.]

Routing Protocols

- If a BGP group is created without any defined peers, a warning message appears when the configuration is committed. [PR/63279: This issue has been resolved.]
- Reverse OIF mappings are lost when you add or delete an interface set of multicast VLANs when subscriber VLANs are active. [PR/423376: This issue has been resolved.]
- When reverse OIF mapping is configured on multicast VLAN interfaces, reverse OIF mappings to DHCP subscriber interfaces are lost if the routing protocol process gracefully restarts. [PR/438930: This issue has been resolved.]
- When the `l3vpn-composite-nexthop` statement and the `multipath vpn-unequal-cost` statement at the `[edit routing-options]` hierarchy level are configured together, the routing process might crash during the multipath calculation for destinations that contain both composite and non-composite eligible paths. [PR/448745: This issue has been resolved.]

Services Applications

- The output of the `show services nat pool` command displays duplicate entries for a single Network Address Translation (NAT) pool. [PR/34678: This issue has been resolved.]

Subscriber Access Management

- Incorrect reverse OIF mappings can be created when a multicast VLAN interface with reverse-OIF mapping enabled receives a join request from a DHCP subscriber and both of the following are true: A valid route to the subscriber is not present

and another route's subnet mask overlaps the address of the subscriber interface. [PR/416774: This issue has been resolved.]

- On MX Series routers, Wimax testing with SBR must be done with Non-Transposable IP for high availability (HA). Otherwise FA-HA authentication will fail with return code 132. [PR/431969: This issue has been resolved.]

VPNs

- On a BGP Layer 3 VPN provider edge router with a combination of (1) label per next hop in the VRFs, (2) configuration of the same IP addresses in different VRFs, and (3) a need for indirect next-hops within the VRFs, then label routes with an indirect next hop might be created incorrectly in the master instance table "mpls.0." [PR/436404: This issue has been resolved.]
- 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 restores traffic to all egress routers. [PR/441392: This issue has been resolved.]
- The VPLS instance on the MX960 router does not learn the remote CE MAC address after the `clear vpls mac-address` command is issued. [PR/476020: This issue has been resolved.]

Related Topics

- New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 6
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 46
- Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers on page 96
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 103

Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers

Changes to the JUNOS Documentation Set

Starting in JUNOS Software Release 9.6, the JUNOS Documentation DVD (DOC-CD-S) will no longer be orderable. The software is still available on the Juniper Networks website and you can download the complete documentation set for later offline viewing or burn your own DVD-ROM by accessing the documentation set at the following URL www.juniper.net/techpubs/resources/cdrom.html.

Errata

This section lists outstanding issues with the documentation.

High Availability

- TX Matrix Plus routers and T1 600 routers that are configured as part of a routing matrix do not currently support nonstop active routing [*Junos High Availability Configuration Guide*].

Interfaces and Chassis

- The `show chassis fabric plane` command is not listed in the Index of Statements and Commands. To access the documentation for this command, see the Table of Contents or use this URL:
https://www.juniper.net/techpubs/en_US/junos9.6/information-products/topic-collections/swcmdref-basics-services/show-chassis-fabric-plane.html.

[*System Basics Command Reference*]

- The following corrections and additions have been made to the descriptions of the command options and output fields for the `show chassis fabric plane` command:

- **[Description]:**

(TX Matrix Plus, T1600, M120, and MX Series routers only) On the M120 router, display the state of all fabric plane connections to the Forwarding Engine Boards (FEBs). On MX Series routers, display the state of all fabric plane connections to the Dense Port Concentrators (DPCs) and Packet Forwarding Engines (PFEs) on the Flexible PIC Concentrators (FPCs). On the TX Matrix Plus router and T1600 routers in a routing matrix, display the state of the fabric management plane and the logical planes on the switch-fabric chassis (SFC) and line-card chassis (LCC). This command can be used on the master Routing Engine only.

- **[Command options]:**

detail—(TX Matrix Plus and T1 600 routers in a routing matrix only) (Optional) Display detailed output for the fabric management plane. Shows Switch Interface Board (SIB) states for the TXP-F13 SIB and TXP-F2S SIB.

extensive—(TX Matrix Plus and T1600 routers in a routing matrix only) (Optional) Display extensive output for the fabric management plane, including the state of the optical links between the F13 SIB on the TX Matrix Plus router and the TXP-T1 600 SIB (ST-SIB-L) on the T1 600 router.

lcc *number*—(TX Matrix Plus router only) (Optional) T1600 router (LCC) that is connected to a TX Matrix Plus router. Replace **number** with a value from 0 through 3.

sfc *number*—(TX Matrix Plus router only) (Optional) Show information about the TX Matrix Plus router (SFC). Replace **number** with 0.

terse—(TX Matrix Plus router only) (Optional) Display terse output for the fabric management plane.

- **[Output field descriptions]:**

Table 2 on page 98 lists the output fields for the **show chassis fabric plane** command. Output fields are listed in the approximate order in which they appear.

Table 2: show chassis fabric plane Output Fields

Field Name	Field Description	Level of output
Plane	(TX Matrix Plus, MX Series, and M120 routers only) Number of the plane.	none
Plane state	(MX Series and M120 routers only) State of each plane: <ul style="list-style-type: none"> ■ ACTIVE—SIB is operational and running. ■ OFFLINE—SIB is powered down. ■ FAULTY— SIB is in alarmed state where the SIB's plane is not operational for the following reasons: <ul style="list-style-type: none"> ■ On-board fabric ASIC is not operational. ■ Fiber optic connector faults. ■ FPC connector faults. ■ SIB mid-plane connector faults. 	none
FEB	(M120 routers only) FEB number and state of links to each FEB: <ul style="list-style-type: none"> ■ Link error—Link between SIB and FPC is not operational. ■ Links ok—Link between SIB and FPC is active. ■ Unused—No FPC is present. 	none
FPC	(MX Series routers only) Slot number of each Dense Port Concentrator (DPC) or Flexible PIC Concentrator (FPC). An FPC occupies two DPC slots on an MX Series router. The interface corresponds to the lowest numbered DPC slot for which the FPC is installed.	none
PFE	(MX Series and M120 routers only) Slot number of each Packet Forwarding Engine and the state of the links to the DCP: Links ok , Link error , or Unused . Each DPC includes four Packet Forwarding Engines. Links ok: Link between SIB and FPC is active. Link error: Link between SIB and FPC is not operational. Unused: No FPC is present.	none

Table 2: show chassis fabric plane Output Fields (continued)

Field Name	Field Description	Level of output
State	<p>(TX Matrix Plus and T1600 routers in a routing matrix only)—State of the fabric plane:</p> <ul style="list-style-type: none"> ■ Online: Fabric plane is operational and running and links on the SIB are operational. ■ Spare: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error. ■ Check: Fabric plane is in alarmed state due to the following reason and the cause of the error must be resolved: <ul style="list-style-type: none"> ■ One or more SIBs (belonging to the fabric plane) in the Online or Spare states has transitioned to the Check state. Check state of the SIB can be caused because of link errors or destination errors. ■ Fault: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the Fault state. A SIB can be in the Fault state because of the following reasons: <ul style="list-style-type: none"> ■ On-board fabric ASIC is not operational. ■ Fiber optic connector faults. ■ FPC connector faults. ■ SIB mid-plane connector faults. ■ Link errors have exceeded the threshold. 	none
Uptime	(TX Matrix Plus and T1600 routers in a routing matrix only)—Time the fabric plane has been up and running.	none
Fabric Management Plane State Output Fields for the show chassis fabric plane extensive command on a TX Matrix Plus router		
PLANE <i>number</i>	<p>State of the fabric plane:</p> <ul style="list-style-type: none"> ■ Online: Fabric plane is operational and running and links on the SIB are operational. ■ Spare: Fabric plane is redundant and can be operational if the operational fabric plane encounters an error. ■ Check: Fabric plane is in alarmed state due to the following reasons and the cause of the error must be resolved: <ul style="list-style-type: none"> ■ One or more SIBs (belonging to the fabric plane) in the Online or Spare states has transitioned to the Check state. Check state of the SIB can be caused because of link errors or destination errors. ■ Fault: Fabric plane is in alarmed state if one or more SIBs belonging to the plane are in the Fault state. A SIB can be in the Fault state because of the following reasons: <ul style="list-style-type: none"> ■ On-board fabric ASIC is not operational. ■ Fiber optic connector faults. ■ FPC connector faults. ■ SIB mid-plane connector faults. ■ Link errors have exceeded the threshold. 	extensive

Table 2: show chassis fabric plane Output Fields (continued)

Field Name	Field Description	Level of output
SIB F13/F2S <i>slot-number</i>	<p>State of the TXP-F13 SIB or TXP-F2S SIB:</p> <ul style="list-style-type: none"> ■ Activating—Transitional state when the SIB is transitioning to the Online or Spare state. ■ Deactivating—Transitional state when the SIB is going offline. ■ Online—SIB is operational and running. ■ Offline—SIB is powered down. ■ Spare—SIB is redundant and will move to active state if one of the working SIBs fail to pass traffic. ■ Empty—No SIB is present. ■ Fault—SIB is in alarmed state due to the following reasons and the cause of the error must be resolved: <ul style="list-style-type: none"> ■ On-board fabric ASIC is not operational. ■ Fiber optic connector faults. ■ FPC connector faults. ■ SIB mid-plane connector faults. ■ Link errors have exceeded the threshold ■ Check—SIB is in alarmed state where the SIB is partially operational due to link or destination errors. Only a SIB that is Online or Spare can transition to the Check state. <p>NOTE: If a SIB is not inserted properly, the SIB cannot transition to the Online or Spare state, and therefore, cannot transition to the Check state.</p>	extensive
SIB F13 <i>slot-number</i> <i>Odd/Even</i>	<p>State of the TXP-F13 SIB Even and Odd port connection optical links from the TX Matrix Plus router (SFC) to the T1600 router (LCC) in the routing matrix. The left four ports on the SFC are labeled Even and provide connections to one even-numbered LCC—LCC0 or LCC2. The right four ports on the SFC are labeled Odd and provide connections to one odd-numbered LCC—LCC1 or LCC3.</p>	extensive
LCC <i>number</i> , SIB <i>slot-number</i>	<p>State of the SIB on the LCC that is connected to the Even or Odd port on the TXP-F13 SIB faceplate:</p> <ul style="list-style-type: none"> ■ Links ok—Links between the TXP-F13 SIB on the SFC and the LCC is active. ■ Link error—Link between the TXP-F13 SIB on the SFC and the LCC is not operational. ■ Unused—No SIB is present. 	extensive
SG <i>number</i> Port <i>number</i>	<p>State of the SG chip ports on the LCC:</p> <ul style="list-style-type: none"> ■ Links ok—Link is active. ■ Link error—Link is not operational. ■ Unused—Port is not in use. 	extensive
SIB F2S <i>slot-number</i>	<p>State of the intra-chassis links between the TXP-F2S and TXP-F13 SIB.</p>	extensive
Fabric Management SIB State Output Fields for the show chassis fabric plane extensive command on a TX Matrix Plus router		

Table 2: show chassis fabric plane Output Fields (continued)

Field Name	Field Description	Level of output
SIB slot-number	<p>State of the SIBs on the T1600 router (LCC) in the routing matrix:</p> <ul style="list-style-type: none"> ■ Activating—Transitional state when the SIB is coming online. ■ Deactivating—Transitional state when the SIB is going offline. ■ Connected—SIBs on an LCC are connected and trained, but are either not online or are spare, because the plane on the the TX Matrix Plus router (SFC) is still offline. The LCC SIB transitions to the Connected state when the F13 SIB to which it connects is online but the SFC plane (to which the LCC SIB connects) is offline for some reason. For instance, when there are insufficient number of F2 SIBs in the plane. ■ Disconnected—If an F13 SIB on the TX Matrix Plus router (SFC) goes offline, then the SIBs on the LCCs connected to the F13 SIB get disconnected. The Disconnected state is valid only for SIBs on an LCC. An LCC SIB transitions to the Disconnected state when the F13 SIB to which it connects goes Offline, irrespective of the state of the SFC plane. <p>NOTE: The Connected and Disconnected states are only applicable to the SIBs on an LCC.</p> <ul style="list-style-type: none"> ■ Online—SIB is operational and running. ■ Offline—SIB is powered down. ■ Spare—SIB is redundant and will move to active state if one of the working SIBs fail to pass traffic. ■ Empty—No SIB is present. ■ Fault—SIB is in alarmed state where the SIB's plane is not operational for the following reasons: <ul style="list-style-type: none"> ■ On-board fabric ASIC is not operational. ■ Fiber optic connector faults. ■ FPC connector faults. ■ SIB mid-plane connector faults. ■ Link errors have exceeded the threshold ■ Check—SIB is in alarmed state where the SIB is partially operational due to link or destination errors. Only a SIB that is Online or Spare can transition to the Check state. <p>NOTE: If a SIB is not inserted properly, the SIB cannot transition to the Online or Spare state, and therefore, cannot transition to the Check state.</p>	extensive
LCC SIB Link State	<p>State of the LCC SIB link:</p> <ul style="list-style-type: none"> ■ Links ok—Link is active. ■ Link error—Link is not operational. ■ Unused—SIB is not in use. 	extensive
SG number Port number	<p>State of the SG chip ports on the LCC:</p> <ul style="list-style-type: none"> ■ Links ok—Link is active. ■ Link error—Link is not operational. ■ Unused—Port is not in use. 	extensive

[*System Basics Command Reference*]

- On M Series, MX Series, and T Series routers, the **targeted-broadcast** statement that is used to forward the direct broadcast packets to the targeted subnet in a network is available in the CLI, but it is not functional for these routers in JUNOS Releases 9.5 to 10.1.

Network Interfaces

- The *Network Interfaces Configuration Guide* in Chapter 5: *Configuring Protocol Family and Interface Address Properties*, the section "Enabling Source Class and Destination Class Usage" contains the following incorrect statement that can be ignored: "On T Series, M120, and M320 routers, the **destination-class** and **source-class** statements are not supported at the [edit firewall family *family-name* filter *filter-name* term *term-name* from] hierarchy level. On other M Series routers, these statements are supported."
- The *Network Interfaces Configuration Guide* in Chapter 5: *Configuring Protocol Family and Interface Address Properties*, the section "Configuring ICCP for MC-LAG" describes functionality not supported for JUNOS Release 9.6 and earlier.
- The PDF version of the *Baseline Network Operations Guide* in Chapter 16: *Use the ping and traceroute Commands*, the section "Check the Accessibility of Two Routers on the Edge" contains an incorrect figure. Figure 17 should show a router network, but instead shows T320 hardware. The correct figure can be located in the HTML version of the section:
http://www.juniper.net/techpubs/en_US/junos9.6/information-products/topic-collections/nog-baseline/ping-traceroute-accessibility-checking.html#id-11574067.

Subscriber Management

The *Subscriber Access Configuration Guide* contains the following dynamic variable errors:

- The *Configuring a Dynamic Profile for Client Access* topic erroneously uses the **\$junos-underlying-interface** variable when configuring an IGMP interface in the client access dynamic profile. The following example provides the appropriate use of the **\$junos-interface-name** variable:

```
[edit dynamic-profiles access-profile]
user@host# set protocols igmp interface $junos-interface-name
```

- Table 25 in the *Dynamic Variables Overview* topic neglects to define the **\$junos-igmp-version** predefined dynamic variable. This variable is defined as follows:

\$junos-igmp-version—IGMP version configured in a client access profile. The JUNOS software obtains this information from the RADIUS server when a subscriber accesses the router. The version is applied to the accessing subscriber when the profile is instantiated. You specify this variable at the [dynamic-profiles *profile-name* protocols igmp] hierarchy level for the **interface** statement.

In addition, the *Subscriber Access Configuration Guide* erroneously specifies the use of a colon (:) when you configure the dynamic profile to define the IGMP

version for client interfaces. The following example provides the appropriate syntax for setting the IGMP interface to obtain the IGMP version from RADIUS:

```
[edit dynamic-profiles access-profile protocols igmp interface $junos-interface-name]
user@host# set version $junos-igmp-version
```

- The *Subscriber Access Configuration Guide* and the *System Basics Configuration Guide* contain information about the `override-nas-information` statement. This statement does not appear in the CLI and is not supported.

User Interface and Configuration

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

VPNs

- In Chapter 10 of the *VPNs Configuration Guide* in the section titled “Configuring BGP Between the PE and CE Routers”, it incorrectly states that route reflectors and cluster IDs are not supported on a routing instance. They are indeed supported, so it is possible to configure the `cluster-id` statement at the `[edit routing-instances routing-instance-name protocols bgp group group-name]` hierarchy level.
- The *Routing Protocols Configuration Guide* and the *VPNs Configuration Guide* both erroneously state that it is not possible to configure route reflectors and cluster IDs for the same routing instance. This type of configuration is now possible. [*Protocols, VPNs*]

Related Topics

- New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 6
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 46
- Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 54
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 103

Upgrade and Downgrade Instructions for JUNOS Release 9.6 for M Series, MX Series, and T Series Routers

This section discusses the following topics:

- Basic Procedure for Upgrading to Release 9.6 on page 104
- Upgrading a Router with Redundant Routing Engines on page 106
- Upgrading the Software for a Routing Matrix on page 106
- Upgrading Using ISSU on page 108

- Upgrading from JUNOS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR on page 108
- Downgrade from Release 9.6 on page 109

Basic Procedure for Upgrading to Release 9.6

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 9.2 you can upgrade to JUNOS Release 9.5 but not to JUNOS Release 9.6. As a workaround, first upgrade to JUNOS Release 9.5 and then upgrade to JUNOS Release 9.6.



NOTE: With JUNOS Release 9.0 and later, the compact flash disk memory requirement for JUNOS Software is 1 GB. For M7i and M10i routers 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) might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the *Junos System Basics Configuration Guide*.

The download and installation process for JUNOS Release 9.6 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-9.6R4.4-domestic-signed.tgz
```

All other customers use the following command:

```
user@host> request system software add validate reboot
source/jinstall-9.6R4.4-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 9.6 **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.



NOTE: Before you upgrade a router that you are using for voice traffic, you should monitor call traffic on each virtual BGF. Confirm that no emergency calls are active. When you have determined that no emergency calls are active, you can wait for non-emergency call traffic to drain as a result of graceful shutdown, or you can force a shutdown. For detailed information on how to monitor call traffic before upgrading, see the *Multiplay Solutions Guide*.

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 the Software for a Routing Matrix

A routing matrix can use either a TX Matrix router as the switch-card chassis (SCC) or a TX Matrix Plus router as the switch-fabric chassis (SFC). By default, when you upgrade software for a TX Matrix router or a TX Matrix Plus router, the new image is loaded onto the TX Matrix or TX Matrix Plus router (specified in the JUNOS CLI by using the **scc** or **sfc** option) and distributed to all T640 routers or T1600 routers in the routing matrix (specified in the JUNOS CLI by using the **lcc** option). To avoid network disruption during the upgrade, ensure the following before beginning the upgrade process:

- A minimum of free disk space and DRAM on each Routing Engine. The software upgrade will fail on any Routing Engine without the required amount of free disk space and DRAM. To determine the amount of disk space currently available on all Routing Engines of the routing matrix, use the CLI `show system storage` command. To determine the amount of DRAM currently available on all the Routing Engines in the routing matrix, use the CLI `show chassis routing-engine` command.
- The master Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and T640 routers or T1600 routers (LCC) are all re0 or are all re1.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and T640 routers or T1600 routers (LCC) are all re1 or are all re0.
- All master Routing Engines in all routers run the same version of software. This is necessary for the routing matrix to operate.
- All master and backup Routing Engines run the same version of software before beginning the upgrade procedure. Different versions of the JUNOS Software can have incompatible message formats especially if you turn on GRES. Because the steps in the process include changing mastership, running the same version of software is recommended.
- For a routing matrix with a TX Matrix router, the same Routing Engine model is used within a TX Matrix router (SCC) and within a T640 router (LCC) of a routing matrix. For example, a routing matrix with an SCC using two RE-A-2000s and an LCC using two RE-1600s is supported. However, an SCC or an LCC with two different Routing Engine models is not supported. We suggest that all Routing Engines are the same model throughout all routers in the routing matrix. To determine the Routing Engine type, use the CLI `show chassis hardware | match routing` command.
- For a routing matrix with a TX Matrix Plus router, the SFC contains two model RE-DUO-C2600-16G Routing Engines, and each LCC contains two model RE-DUO-C1800-8G Routing Engines.



NOTE: It is considered best practice to make sure that all master Routing Engines are re0 and all backup Routing Engines are re1 (or vice versa). For the purposes of this document, the master Routing Engine is re0 and the backup Routing Engine is re1.

To upgrade the software for a routing matrix, perform the following steps:

1. Disable graceful Routing Engine switchover (GRES) on the master Routing Engine (re0) and save the configuration change to both Routing Engines.
2. Install the new JUNOS Software release on the backup Routing Engine (re1) while keeping the currently running software version on the master Routing Engine (re0).
3. Load the new JUNOS Software on the backup Routing Engine. After making sure that the new software version is running correctly on the backup Routing Engine (re1), switch mastership back to the original master Routing Engine (re0) to activate the new software.
4. Install the new software on the new backup Routing Engine (re0).

For the detailed procedure, see the [Routing Matrix with a TX Matrix Feature Guide](#) or the [Routing Matrix with a TX Matrix Plus Feature Guide](#).

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 High Availability Configuration Guide*.

Upgrading from JUNOS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR

JUNOS Release 9.3 introduced NSR support for PIM for IPv4 traffic. However, the following PIM features are not currently supported with 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 9.3 introduced 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 if one of the unsupported PIM features is enabled but NSR is not enabled, no additional steps are necessary and you can use the standard upgrade procedure described in other sections of these instructions. If NSR is enabled and no NSR-incompatible PIM features are enabled, use the standard reboot or ISSU procedures described in the other sections of these instructions.

Because the **nonstop-routing disable** statement was not available in JUNOS Release 9.2 and earlier, if both NSR and an incompatible PIM feature are enabled on a router to be upgraded from JUNOS Release 9.2 or earlier to a later release, you must disable PIM before the upgrade and reenabling it after the router is running the upgraded JUNOS Software and you have entered the **nonstop-routing disable** statement. If your router is running JUNOS Release 9.3 or later, you can upgrade to a later release without disabling NSR or PIM—simply use the standard reboot or ISSU procedures described in the other sections of these instructions.

To disable and reenabling 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 JUNOS Release 9.3 or later 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 the upgraded JUNOS Software, enter configuration mode, disable PIM NSR with the `nonstop-routing disable` statement, and then reenabling PIM:

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

Downgrade from Release 9.6

To downgrade from Release 9.6 to another supported release, follow the procedure for upgrading, but replace the 9.6 `install` 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**
- New Features in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 6
 - Changes in Default Behavior and Syntax in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 46
 - Issues in JUNOS Release 9.6 for M Series, MX Series, and T Series Routers on page 54
 - Errata and Changes in Documentation for JUNOS Software Release 9.6 for M Series, MX Series, and T Series Routers on page 96

JUNOS Software Release Notes for Juniper Networks SRX Series Services Gateways

Powered by JUNOS Software, the Juniper Networks SRX Series Services Gateways provide robust networking and security services. SRX Series Services Gateways range from lower-end devices designed to secure small distributed enterprise locations to high-end devices designed to secure enterprise infrastructure, data centers, and server farms. This series of devices includes SRX100, SRX210, SRX240, SRX650, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

- New Features in JUNOS Release 9.6 for SRX Series Services Gateways on page 110
- Changes In Default Behavior and Syntax in JUNOS Release 9.6 for SRX Series Services Gateways on page 123
- Known Limitations in JUNOS Release 9.6 for SRX Series Services Gateways on page 125
- Issues in JUNOS Release 9.6 for SRX Series Services Gateways on page 130
- Errata and Changes in Documentation for JUNOS Release 9.6 for SRX Series Services Gateways on page 143

New Features in JUNOS Release 9.6 for SRX Series Services Gateways

The following features have been added to JUNOS Release 9.6. Following the description is the title of the manual or manuals to consult for further information.

- Software Features on page 111
- Hardware Features—SRX100 Services Gateway on page 119
- Hardware Features—SRX210 Services Gateways on page 122
- Hardware Features—SRX5600 and SRX5800 Services Gateways on page 122

Software Features

Chassis Clustering

- **Redundancy group IP address monitoring**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

SRX3000 and SRX5000 line devices in a chassis cluster now include IP address monitoring for redundancy groups. IP monitoring works as a ping to upstream network devices trying to determine whether or not a given part of the network infrastructure is reachable and then deducts the configured failover value from the threshold priority if it is not. Once the configured failover threshold for IP address monitoring is reached, the failover value for IP monitoring is deducted from the redundancy group threshold value.

The IP address monitoring value, along with interface monitoring, SPU monitoring, cold-sync monitoring, and—for the SRX3000 line only—NPC monitoring values are used to determine whether or not a redundancy group will fail over. The primary IP addresses that should be monitored are the device addresses to ensure that valid traffic coming into the device can be forwarded to the appropriate network device.

The interval to check reachability is once per second. After failing to reach a configured IP address for five consecutive attempts, the IP is determined to be unreachable and the failover value is deducted from the redundancy group's priority threshold. The failover value will be reinstated if the IP address becomes reachable again.

The maximum number of monitoring IPs that can be configured per cluster is 32 for the SRX3000 line and 64 for the SRX5000 line.

Monitoring can be accomplished only if the IP address is reachable on a redundant Ethernet (reth) interface, and IP addresses cannot be monitored over a tunnel. The feature also cannot be used on a cluster running in transparent mode.

[Junos OS Security Configuration Guide]

- **Low-impact ISSU chassis cluster upgrades**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

The SRX3000 and SRX5000 lines now support in-service software upgrades (ISSUs) to eliminate downtime when you are upgrading devices in a chassis cluster. You can now run a single request command to upgrade both devices in an SRX3000 or SRX5000 line cluster without needing to either reset devices or take the cluster out of service. This feature is supported only on JUNOS Release 9.6 and later, so when you initially upgrade to Release 9.6, you will still need to upgrade devices in a cluster separately.

[Junos OS Security Configuration Guide]

- **Active/active chassis clustering**—This feature is now supported on SRX240 and SRX650 devices in addition to existing support on SRX210, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

The data plane now supports active/active chassis clustering for these SRX Series devices. The chassis clustering on these SRX Series devices is no longer restricted to the creation of only one redundancy group beyond redundancy group 0. You

can now configure one or more redundancy groups numbered 1 through 128. Multiple redundancy groups make it possible for traffic to arrive on an interface of one redundancy group and egress on an interface that belongs to another redundancy group. In this situation, the ingress and egress interfaces might not be active on the same node. When this happens, the traffic is forwarded over the fabric link to the appropriate node. SRX Series chassis clusters operate with an active/backup control plane.

[Junos OS Security Configuration Guide]

- **Active/passive chassis clustering**—This feature is now supported on SRX240 and SRX650 devices in addition to existing support on SRX210, SRX3400, SRX3600, SRX5600, and SRX5800 devices.

In an active/passive chassis cluster, a single device in the cluster is used to route all traffic while the other device keeps session state and configuration synchronization with the active and is used only in the event of a failure. When a failure occurs, the backup device becomes master and controls all forwarding. An active/passive chassis cluster can be achieved by using redundant Ethernet interfaces (reths) that are all assigned to the same redundancy group. If any of the interfaces in an active group in a node fails, the group is declared inactive and all the interfaces in the group will fail over to the other node.

[Junos OS Security Configuration Guide]

- **Control link recovery**—This feature is now supported on SRX210, SRX240, and SRX650 devices in addition to existing support on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

Previously, when a node was disabled due to control link failure, after you fixed the issue, you had to manually reboot the disabled node allow it to rejoin the cluster. With JUNOS Release 9.6, you can specify that control link recovery be done automatically by the system by using the **set chassis cluster control-link-recovery** command (this feature is disabled by default). Once the system determines that the control link is healthy, it issues an automatic reboot on the disabled node. When the disabled node reboots, the node rejoins the cluster. There is no need for any manual intervention.

[Junos OS Security Configuration Guide]

- **Cold synchronization monitoring**—This feature is now supported on SRX240 and SRX650 devices in addition to existing support on all other SRX Series devices.

The process of synchronizing data plane runtime objects (RTOs) on the startup of the Services Processing Units (SPUs) or flowd is called cold sync. Chassis clustering supports the process of monitoring the cold-sync state of all SPUs or flowd on a node. Also, if you enable preempt, cold-sync monitoring prevents the node from taking over mastership until the cold-sync process is completed for all the SPUs or flowd on the node.

[Junos OS Security Configuration Guide]

- **SNMP failover traps**—This feature is now supported on SRX240 and SRX650 devices in addition to existing support on all other SRX Series devices.

Chassis clustering supports SNMP traps, which are triggered whenever there is a redundancy group failover. You can specify that a trace log be generated by using the **set chassis cluster traceoptions flag snmp** command.

[Junos OS Security Configuration Guide]

- **Dial-up VPNs**—This feature is now supported on the SRX3000 and SRX5000 lines of devices in addition to existing support on SRX210, SRX240, and SRX650 devices.

Dial-up VPNs are now supported in chassis clustering environments.

[Junos OS Security Configuration Guide]

Flow and Processing

- **Selective stateless packet-based services**—This feature is supported on SRX100, SRX210, SRX240, and SRX650 devices.

Selective stateless packet-based services allow you to use both flow-based and packet-based forwarding simultaneously on a system. You configure these services by specifying an action modifier in stateless firewall filters and applying the filters to certain interfaces. On these interfaces, traffic that matches the filter criteria will bypass flow-based forwarding. Bypassing flow-based forwarding can be useful for traffic for which you explicitly want to avoid session-scaling constraints.

A defined set of stateless services is available with selective stateless packet-based services.

[Junos OS Administration Guide for Security Devices]

- **Datapath debugging**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

In datapath debugging, a packet goes through multiple Services Processing Units (SPUs). At the same time, several FPC I/O cards (IOCs) provide EZchip ingress and egress traffic management. JUNOS Software supports datapath debugging for filter-based, per-packet counting and logging to record the processing path of a packet. Only the matched packets are traced by the IOC EZchip ingress, EZchip egress, load-balancing thread (LBT), and packet-ordering thread (POT).

An SPU dedicated to central point functionality is called a large central point. However, when such a dedicated SPU is not affordable, you can configure one SPU to perform both normal flow processing and the functions of the central point. The SPU and the central point share the same LBT and POT infrastructure.

The following events are defined for the SRX5600 and SRX5800 devices in the packet processing path:

- ezchip.ingress
- ezchip.egress
- spu.lbt
- spu.pot

[Junos OS Security Configuration Guide]

- **Network processor bundling**—This feature is supported on SRX5600 and SRX5800 devices.

This feature enables distribution of data traffic from one interface to multiple network processors for packet processing. A primary network processor is assigned for an interface that receives the ingress traffic and distributes the packets to several other secondary network processors. A single network processor can act as a primary network processor or a secondary network processor to multiple interfaces. A single network processor can join only one network processor bundle.

[Junos OS Security Configuration Guide]

Interfaces and Routing

- **Layer 2 transparent mode chassis clusters**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

A pair of SRX3400, SRX3600, SRX5600, or SRX5800 Services Gateways in Layer 2 transparent mode can be connected in a chassis cluster to provide network node redundancy. When configured in a chassis cluster, one node acts as the primary device and the other as the secondary device, ensuring stateful failover of processes and services in the event of system or hardware failure. If the primary device fails, the secondary device takes over processing of traffic. If the primary device fails, the physical ports in the failed device become inactive (go down) for a few seconds before they become active (come up) again.



NOTE: In JUNOS Release 9.6, devices in Layer 2 transparent mode can be deployed only in active/passive chassis cluster configurations.

On a device in a Layer 2 transparent mode chassis cluster, the redundant Ethernet interface is configured as a Layer 2 logical interface. Physical interfaces are bound to the parent redundant Ethernet interface. To configure a redundant Ethernet (reth) interface as a Layer 2 logical interface, use the **family bridge** statement at the [edit interfaces] hierarchy level. The redundant Ethernet interface can be configured as either an access interface or a trunk interface.

All JUNOS screen options that are available for a single, nonclustered device are available for devices in Layer 2 transparent mode chassis clusters.

[Junos OS Interfaces and Routing Configuration Guide]

- **Layer 2 transparent mode VLAN retagging**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

The VLAN identifier in packets arriving on a Layer 2 trunk port can be rewritten or “retagged” with a different internal VLAN identifier. For example, an incoming packet with the VLAN identifier 11 can be retagged with the VLAN identifier 2. VLAN retagging is a symmetric operation; a packet with the VLAN identifier 2 is retagged with VLAN identifier 11 when it exits the same trunk port. The trunk port can be a redundant Ethernet interface in a Layer 2 transparent mode chassis cluster configuration.

To configure VLAN retagging on a Layer 2 trunk port, use the **vlan-rewrite translate** statements at the [edit interfaces] hierarchy level. Specify the VLAN identifier of the incoming packets that are to be retagged and the internal VLAN identifier to be written in the packets. The VLAN identifier of the incoming packets must not

be the same VLAN identifier configured with the **native-vlan-id** statement for the trunk port. The internal VLAN identifier must be in the VLAN identifier list for the trunk port.

[Junos OS Interfaces and Routing Configuration Guide]

Intrusion Detection and Prevention (IDP)

- **Performance and capacity tuning for IDP**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

If you are deploying IDP policies, you can tune the device to increase IDP session capacity. By using the provided commands to change the way the system allocates resources, you can achieve a higher IDP session capacity. To configure IDP dedicated mode, use the **maximize-idp-sessions weight idp** statement at the **[set security forwarding-process application-services]** hierarchy level.

[Junos OS Security Configuration Guide]

- **IDP multiple detector support**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices. These devices now support loading multiple IDP detector simultaneously. When a policy is loaded, it is also associated with a detector. If the new policy being loaded has an associated detector that matches the detector already being used by the existing policy, the new detector is not loaded and both policies use a single associated detector. But if the new detector does not match the current detector, the new detector is loaded along with the new policy. In this case, each loaded policy will then use its own associated detector for attack detection. Note that a maximum of two detectors can be loaded at any given time.

[Junos OS Security Configuration Guide]

IPsec VPN

- **Hub-and-spoke VPNs**—This feature is supported on all SRX Series devices.

For route-based IPsec VPNs, JUNOS Software supports a topology called hub-and-spoke, which allows two or more remote sites to communicate to each other through a central site using VPN tunnels. To use hub-and-spoke, you create two IPsec VPN tunnels that terminate at a gateway and define a pair of routes so that the device directs traffic exiting from one tunnel to the other tunnel. Only route-based hub-and-spoke is supported.

[Junos OS Security Configuration Guide]

J-Web

- **J-Web user interface enhancements**—This feature is supported on all SRX Series devices.

The J-Web user interface has been significantly updated in JUNOS Release 9.6. The menu system has been updated to provide more intuitive navigation to most pages. For example, configuration pages for firewall policies and UTM policies are now grouped together under the **Configure > Security > Policy** menu. In addition, the look and feel of the monitoring and configuration pages has been updated in JUNOS Release 9.6.

[Junos OS Security Configuration Guide, Junos OS Interfaces and Routing Configuration Guide, Junos OS Administration Guide for Security Devices]

Management and administration

- **Redundant system log server**—This feature is supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices.

For purposes of logging redundancy, it is now possible to configure two external system log servers for SRX3000 and SRX5000 line devices. The feature is supported on both standalone and cluster-paired devices. Note that, while the feature will work on both active/backup and active/active chassis clusters, an active/active pair can already be configured to allow each node to send to a separate system log server. This feature is primarily of use to standalone and active/backup cluster deployments. Note also that configuring a second system log server might result in performance degradation because of increased logging processing and traffic.

[Junos OS Administration Guide for Security Devices]

- **Support for the TFTPBOOT installation method**—This feature is supported on SRX100 devices in addition to existing support on SRX210 and SRX650 devices.

You install the JUNOS Software by using the Trivial File Transfer Protocol BOOT (TFTPBOOT) method. During installation of the JUNOS Software, the secondary boot loader in the services gateway retrieves the JUNOS Software package from a TFTP server. The software image is then installed on the internal flash. Using TFTP installation to install a new image will wipe out any user-generated

configurations on the device. The device will come up with the factory default configuration.



NOTE: The TFTPBOOT method can be used only on LANs.

To install the software image on the internal flash, use the following command at the loader prompt:

Loader > install *URL*

where *URL* is *tftp://tftp server ip package name*

You can use the TFTPBOOT method in the following scenarios:

- To bring up the SRX100 Services Gateway if the standard boot process fails
- To install the JUNOS Software on the SRX100 Services Gateway for the first time
- To start JUNOS without using the NAND flash

[Junos OS Administration Guide for Security Devices]

- **NAT information in session logs**—This feature is supported on all SRX Series devices.

As with previous releases, a session can be logged whenever it is created, closed, rejected, or denied. Starting with JUNOS Release 9.6, the following additional information is included in session logs for all SRX Series devices:

- **nat-source-address**—The translated NAT source address if NAT was applied; otherwise, the source-address.
- **nat-source-port**—The translated NAT source port if NAT was applied; otherwise, the source-port.
- **nat-destination-address**—The translated NAT destination address if NAT was applied; otherwise, the destination-address.
- **nat-destination-port**—The translated NAT destination port (if any); otherwise, the destination-port.
- **dst-nat-rule-name**—The destination NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if destination address translation takes place, then this field shows the static NAT rule name.
- **src-nat-rule-name**—The source NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if source address translation takes place, then this field shows the static NAT rule name.



NOTE: Some sessions might have both destination and source NAT applied and the information logged.

- **service-name**—The service (application) through which the packet traversed.
- **session-id-32**—The 32-bit session ID

Both traditional and structured system log formats are supported. The system logs can be exported both to the Routing Engine (RE) and to an external host.

[Junos OS Security Configuration Guide]

Security

■ Network Address Translation (NAT)

- **Disabling port randomization for source NAT**—This feature is supported on all SRX Series devices.

For pool-based source NAT and interface NAT, port numbers are allocated randomly by default. While randomized port number allocation can provide protection from security threats, such as DNS poison attacks, it can also affect performance and memory usage for pool-based source NAT.

You can disable port randomization by using the **port-randomization disable** statement at the **[edit security nat source]** hierarchy level. To re-enable port randomization, use the **port-randomization** statement at the **[edit security nat source]** hierarchy level.

[Junos OS Security Configuration Guide]

- **Persistent NAT**—This feature is supported on SRX100, SRX210, SRX240, and SRX650 devices.

Configuration of persistent NAT allows applications to use the Session Traversal Utilities for NAT (STUN) protocol for NAT traversal.

The STUN protocol is a client-server protocol that allows a client application behind a NAT device to learn its public IP address and the type of NAT used to allocate the address bindings. The STUN client is commonly used with Session Initiation Protocol (SIP) VoIP applications where IP addresses and port numbers are encoded within the application data. Both the STUN client and STUN server are provided by the application.

Persistent NAT ensures that all requests from the same internal transport address are mapped to the same external transport address by the NAT device closest to the STUN server. The following types of persistent NAT can be configured on the Juniper Networks device:

- **Any remote host**—All requests from a specific internal IP address and port are mapped to the same external IP address and port. Any external host can send a packet to the internal host using the mapped external address when the incoming policy from external to internal is configured.
- **Target host**—All requests from a specific internal IP address and port are mapped to the same external IP address and port. An external host

can send a packet to an internal host only if the internal host had previously sent a packet to the external host's IP address.

- **Target host port**—All requests from a specific internal IP address and port are mapped to the same external IP address and port. An external host can send a packet to an internal host only if the internal host had previously sent a packet to the external host's IP address and port.

To configure persistent NAT options, use the **persistent-nat** statement in the [edit security nat] hierarchy level.

You can configure security policies with two new predefined services, **junos-stun** and **junos-persistent-nat**, to permit or deny persistent NAT traffic.

[Junos OS Security Configuration Guide]



NOTE: Persistent NAT is sometimes referred to as Cone NAT. The term Cone NAT has been replaced by Persistent NAT by the IETF.

■ UTM

- **Unified Threat Management J-Web support**— The UTM J-Web Quick Configuration screens have been redesigned to support a new J-Web framework for enhanced usability.

[Junos OS Security Configuration Guide]

Hardware Features—SRX100 Services Gateway

JUNOS Software for the SRX100 Services Gateway integrates the world-class network security and routing capabilities of Juniper Networks. JUNOS Software for the SRX100 Services Gateway includes a wide range of security services, including policies, screens, Network Address Translation (NAT), and other packet-based services, which are also supported on the other SRX Series Services Gateways.

The SRX100 Services Gateway offers features that provide complete functionality and flexibility for delivering secure Internet and intranet access. This services gateway offers stable, reliable, and efficient IP routing in addition to switching support and LAN connectivity. The services gateway provides Internet Protocol Security (IPsec), VPN, and firewall services for small and medium companies and enterprise branch and remote offices.

The SRX100 device can be connected directly to traditional private networks, such as leased line, Frame Relay, and MPLS networks, or the public Internet.

There are two variants of the SRX100 device:

- Low Memory
- High Memory

The SRX100 device has redundant and resilient hardware. Table 3 on page 121 provides the SRX100 device chassis specifications.

Table 3: SRX100 Services Gateway Chassis Specifications

Description	Value
Chassis height	1.38 in. (35 mm)
Chassis width	8.5 in. (216 mm)
Chassis depth	5.79 in. (147 mm)
Chassis weight	1.86 lb (844 g)
Altitude	No performance degradation to 10,000 ft (3048 m)
Temperature	<p>Normal operation ensured in temperature range of 32°F (0°C) to 104°F (+ 40°C)</p> <p>Nonoperating storage temperature in shipping container: -40°F (-40°C) to 158°F (70°C)</p> <p>NOTE: The SRX100 Services Gateway operating temperature is 35°C when installed in a rack.</p>
Maximum thermal output	<p>The maximum thermal values for the two types of services gateways are as follows:</p> <ul style="list-style-type: none"> ■ Low Memory – AC power: 73 BTU/hour (21.5 W) ■ High Memory – AC power: 73 BTU/hour (21.5 W) <p>NOTE: These specifications are estimates and subject to change.</p>
Power supply adapter	30 watts

Table 4 on page 121 provides information about the SRX100 device hardware features.

Table 4: SRX100 Services Gateway Hardware Features

Feature	Description
Fast Ethernet	Eight ports on the front panel provide LAN connectivity to hubs, switches, local servers, and workstations with link speeds of 10/100 Mbps.
Universal serial bus	One port on the front panel supports a USB storage device that can function as a secondary boot device in the event of internal flash failure. The USB port also provides interfaces for communicating with peripherals such as USB storage devices and USB storage-device adapters.
Console	One port on the front panel functions as a management port for directly logging in to a device to configure it by using the CLI.
External power supply	<p>The total power consumption by the two SRX100 Services Gateway variants is as follows:</p> <ul style="list-style-type: none"> ■ Low Memory—20 W @ 12 V ■ High Memory—22 W @ 12 V

Table 4: SRX100 Services Gateway Hardware Features *(continued)*

Memory	<ul style="list-style-type: none"> ■ Fixed Random Access: <ul style="list-style-type: none"> ■ Low Memory—512 MB ■ High Memory—1 GB ■ NAND flash—4 MB ■ Internal flash—1 GB
--------	---

[*SRX100 Services Gateway Hardware Guide*]

Hardware Features—SRX210 Services Gateways

Support for 3G Wireless Functionality—JUNOS Release 9.6 supports 3G wireless functionality on SRX210 devices to provide high-speed wireless WAN connectivity for both primary and backup devices.

Third-generation (3G) networks are wide area cellular telephone networks that have evolved to include high-data rate services of up to 3 Mbps.

SRX210 devices have an ExpressCard slot on the back panel.

Following Juniper Wireless Modems are supported in the SRX210 device in JUNOS 9.6:

- EXPCD-3G-CDMA-V: 3G EVDO ExpressCard for Verizon Wireless, available from Juniper starting October, 2009.
- EXPCD-3G-CDMA-S: 3G EVDO ExpressCard for Sprint, , available from Juniper starting November, 2009.

In addition, SRX210 devices running JUNOS 9.5 or JUNOS 9.6, supports following Sierra AirCards:

- Sierra Wireless AirCard Global System for Mobile Communications (GSM) High-Speed Downlink Packet Access (HSDPA) ExpressCard - Sierra Wireless AirCard 880E

For more information on installing 3G ExpressCards, see the *SRX210 Services Gateway Hardware Guide*.

For more information on configuring the 3G interface, see the *Junos OS Interfaces and Routing Configuration Guide*.

Hardware Features—SRX5600 and SRX5800 Services Gateways

16x1GE-SFP port module—JUNOS Release 9.6 supports the new 16x1GE-SFP port module for the Flex I/O Card (Flex IOC) of the SRX5600 and SRX5800 Services Gateways. This port module has 16 1-gigabit SFP Ethernet ports.

[*Junos OS Interfaces and Routing Configuration Guide*]

- Related Topics**
- Known Limitations in JUNOS Release 9.6 for SRX Series Services Gateways on page 125
 - Issues in JUNOS Release 9.6 for SRX Series Services Gateways on page 130
 - Errata and Changes in Documentation for JUNOS Release 9.6 for SRX Series Services Gateways on page 143

Changes In Default Behavior and Syntax in JUNOS Release 9.6 for SRX Series Services Gateways

Chassis Cluster

- On SRX650 devices in chassis cluster mode, the CT1/E1 PIC goes offline and does not come online.

CLI

- On SRX100, SRX210, and SRX650 devices, the `show security monitoring` command and associated SNMP are not available.
- On SRX5600 and SRX5800 devices, the `set security end-to-end-debug` CLI hierarchy has been changed to `set security datapath-debug`.

DHCP

- On SRX Series devices, the DHCP client is supported on ATM interfaces (ADSL, G.SHDSL) along with existing support on Ethernet interfaces.

Flow and Processing

- On SRX650 devices, although the physical installed DRAM is 2 GB and uboot detected is 2 GB, JUNOS Software detects only 1GB.
- On SRX Series devices, the factory default for the maximum number of backup configurations allowed is 5. Therefore, you can have one active configuration and a maximum of five rollback configurations. Increasing this backup configuration number will result in increased memory usage on disk and increased commit time.

To modify the factory defaults, use the following commands:

```
root@host# set system max-configurations-on-flash number
```

```
root@host# set system max-configuration-rollbacks number
```

where `max-configurations-on-flash` indicates backup configurations to be stored in the configuration partition and `max-configuration-rollbacks` indicates the maximum number of backup configurations.

Interfaces and Routing

- To minimize the size of system logs, the default logging level in the factory configuration has changed from **any** to **any critical**.
- On the SRX3000 and SRX5000 line devices, the **set protocols bgp family inet flow** and **set routing-options flow** CLI statements are no longer available because BGP flow spec functionality is not supported on these platforms.

Intrusion Detection and Prevention (IDP)

- With compressed DFA, the application signature will have a different filename `/var/db/idpd/bins/compressed_ai.bin`, instead of the current `/var/db/idpd/bins/compiled_ai.bin`.
- While running commands in IDP, ensure that you provide the service field values for custom attack definitions in lowercase.

In the following example, the protocol service field value **udp** is specified in lowercase:

```
set security idp custom-attack temp severity info attack-type signature context packet
direction any pattern .* protocol udp destination-port match equal value 1333
```

- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, for brute force and time binding related attacks, the logging is to be done only when match **count** is equal to the **threshold**. It means there could be only one log generated within a 60 seconds time period in which the threshold is measured. It will stop any repetitive log that would have been generated otherwise. It will also keep the behavior consistent with other IDP platforms like IDP-Standalone.
- The IDP **ip-action** statement is now supported on TCP, UDP, and ICMP flows. When the ip-action target is service, the ip-action flow is applied if the traffic matches the values specified for source port, destination port, source address and destination address. However, for ICMP flows, the destination port is 0 so that any ICMP flow matching source port, source address, and destination address would be blocked. For more information, see the *Junos OS CLI Reference*.

J-Web

- On SRX3400, SRX3600, SRX5600, SRX5800 devices, to add the Predefined Attacks and Predefined Attack Groups the user need not type the attack names. The Predefined Attacks and Predefined Attack Groups lists are available. Select the attacks from the list and click the left arrow to add them.
- For SRX100, SRX210, SRX240, and SRX650 devices, the LED status (Alarm, HA, ExpressCard, Power Status, and Power) shown in the front panel for Chassis View does not replicate the exact status of the device.

Management and Administration

- The following session logging fields have been renamed in JUNOS Release 9.6:

Old Name	New Name
inbound-packets	packets-from-client
inbound-bytes	bytes-from-client
outbound-packets	packets-from-server
outbound-bytes	bytes-from-server

- Previously, security logs were always timestamped using the UTC time zone. In JUNOS Release 9.6, you can use the **set system time-zone** CLI command to specify the local time zone that the system should use to timestamp the security logs. If you want to timestamp logs using the UTC time zone, use the **set system time-zone utc** and **set security log utc-timestamp** CLI statements.

Known Limitations in JUNOS Release 9.6 for SRX Series Services Gateways

[accounting-options] Hierarchy

- For SRX210 and SRX240 devices, the **accounting**, **source-class**, and **destination-class** statements in the [accounting-options] hierarchy level are not supported.

Chassis Cluster

- Multicast traffic streams are not supported on SRX Series chassis clusters.
- For SRX3000 and SRX5000 line chassis clusters, screen statistics data can be gathered on the primary device only.
- The IDP feature is not supported in active/active chassis clustering.

Command-Line Interface (CLI)

On SRX210 and SRX240 devices, J-Web crashes if more than nine users log in to the device using the CLI.

The number of users allowed to access the device is limited as follows:

- For SRX210 devices: four CLI users and three J-Web users
- For SRX240 devices: six CLI users and five J-Web users

Flow and Processing

Maximum concurrent SSH, Telnet, and Web sessions—For SSH, Telnet, and Web sessions, the maximum number of concurrent sessions is as follows:

Sessions	SRX210	SRX240	SRX650
ssh	3	5	5
telnet	3	5	5
Web	3	5	5



NOTE: These defaults are provided for performance reasons.

Hardware

- **Transceivers**—We strongly recommend that only Juniper Networks transceivers be used on SRX Series and J Series interface modules. Different transceiver types (long-range, short-range, copper, and others) can be used together on multiport SFP interface modules as long as they are provided by Juniper Networks. We cannot guarantee that the interface module will operate correctly if third-party transceivers are used.

Please contact Juniper Networks for the correct transceiver part number for your device.

- **Filters and policing**—This section covers the filter and policing limitations:
On SRX3400 and SRX3600 devices, the following feature is not supported by a simple filter:

- Forwarding class as match condition

On SRX3400 and SRX3600 devices, the following features are not supported by a policer or a three-color-policer:

- Color-aware mode of a three-color-policer
- Filter-specific policer
- Forwarding class as action of a policer
- Logical interface policer
- Logical interface three-color policer
- Logical interface bandwidth policer

- Packet loss priority as action of a policer
- Packet loss priority as action of a three-color-policer

On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the following features are not supported by a firewall filter:

- Policer action
- Egress FBF
- FTF

SRX3400 and SRX3600 devices have the following limitations of a simple filter:

- In the packet processor on an IOC, up to 100 logical interfaces can be applied with simple filters.
- In the packet processor on an IOC, the maximum number of terms of all simple filters is 4000.
- In the packet processor on an IOC, the maximum number of policers is 4000.
- In the packet processor on an IOC, the maximum number of three-color-policers is 2000.
- The maximum burst size of a policer or three-color-policer is 16 MB.

IGMP

- SRX100 devices do not support IGMP snooping.

Interfaces and Routing

- On SRX650 devices, MAC pause frame and FCS error frame counters are not supported for the interfaces `ge-0/0/0` through `ge-0/0/3`.
- On SRX240 devices, IP multicast switching is not supported; because of this, multicast snooping is based on corresponding IP multicast Layer 2 address (01:00:5e:xx:xx:xx). On SRX240 devices, all multicast receivers with an IP multicast address mapped to the same Layer 2 address will receive the packets.
- On SRX240 and SRX650 devices, the VLAN range from 3967 to 4094 falls under the reserved VLAN address range, and the user is not allowed any configured VLANs from this range.
- On SRX650 devices, the last 4 ports of a 24-Gigabit Ethernet switch GPIM can be used either as RJ-45 or SFP ports. If both are present and providing power, the SFP media is preferred. If the SFP media is removed or the link is brought down, then the interface will switch to the RJ-45 medium. This can take up to 15 seconds, during which the LED for the RJ-45 port might go up and down intermittently. Similarly when the RJ-45 medium is active and an SFP link is brought up, the interface will transition to the SFP medium, and this transition could also take a few seconds.

- The user can use IPsec only on an interface that resides in the routing instance `inet 0`. The user will not be able to assign an internal or external interface to the IKE policy if that interface is placed in a routing instance other than `inet 0`.
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the following multicast IPv6 and MVPN CLI commands are not supported. However, if you enter these commands in the CLI editor, they will appear to succeed and will not display an error message.
 - `show pim interfaces inet6`
 - `show pim neighbors inet6`
 - `show pim source inet6`
 - `show pim rps inet6`
 - `show pim join inet6`
 - `show pim mvpn`
 - `show multicast next-hops inet6`
 - `show multicast rpf inet6`
 - `show multicast route inet6`
 - `show multicast scope inet6`
 - `show multicast pim-to-mld-proxy`
 - `show multicast statistics inet6`
 - `show multicast usage inet6`
 - `show msdp sa group group`
 - `set protocols pim interface interface family inet6`
 - `set protocols pim disable interface interface family inet6`
 - `set protocols pim family inet6`
 - `set protocols pim disable family inet6`
 - `set protocols pim apply-groups group disable family inet6`
 - `set protocols pim apply-groups group family inet6`
 - `set protocols pim apply-groups-except group disable family inet6`
 - `set protocols pim apply-groups group interface interface family inet6`
 - `set protocols pim apply-groups group apply-groups-except group family inet6`
 - `set protocols pim apply-groups group apply-groups-except group disable family inet6`
 - `set protocols pim assert-timeout timeout-value family inet6`

- `set protocols pim disable apply-groups group family inet6`
- `set protocols pim disable apply-groups-except group family inet6`
- `set protocols pim disable export export-join-policy family inet6`
- `set protocols pim disable dr-election-on-p2p family inet6`
- `set protocols pim dr-election-on-p2p family inet6`
- `set protocols pim export export-join-policy family inet6`
- `set protocols pim import export-join-policy family inet6`
- `set protocols pim disable import export-join-policy family inet6`
- On SRX100, SRX210, SRX240, and SRX650 devices, flow mode does not support asymmetric routing for stateful sessions. As a result of this behavior, trace-route might not work when VRRP is configured across SRX Series devices.

Intrusion Detection and Prevention (IDP)

- On SRX Series devices, IP actions do not work when you select a timeout value greater than 65535 in the IDP policy.
- The maximum number of IDP sessions supported is as follows:
 - SRX210–16,000
 - SRX240–32,000
 - SRX650–128,000
- JUNOS Release 9.6 supports all IDP policy templates except All Attacks. There is a 100-MB policy size limit for integrated mode and a 150-MB policy size limit for dedicated mode, and the current IDP policy templates supported are dynamic, based on the attack signatures being added. Therefore, be aware that supported templates might eventually grow past the policy-size limit.

The following IDP policies are supported on SRX Series devices:

- DMZ_Services
- DNS_Service
- File_Server
- Getting_Started
- IDP_Default
- Recommended
- Web_Server
- By default, the detector embedded in the SRX Series devices has the SIP, SSL, SSH, and MSPRC protocol decoders enabled.

- IDP failover is not supported in chassis clustering.
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, an address book entry with a dns-name option resolves an invalid FQDN to a network address, which is not a valid host address. [PR/465881]

NetScreen-Remote

- NetScreen-Remote is not supported on SRX Series devices.

System

- By default, the detector embedded in the SRX Series devices has the SIP, SSL, SSH, and MSPRC protocol decoders disabled.
- On the four Gigabit Ethernet ports (**ge-0/0/0** through **ge-0/0/3**) of an SRX650 device, if a port is linked up at 10 or 100 Mbps, it will not support jumbo frames. Frames greater than 1500 bytes are dropped.

- Related Topics**
- New Features in JUNOS Release 9.6 for SRX Series Services Gateways on page 110
 - Issues in JUNOS Release 9.6 for SRX Series Services Gateways on page 130
 - Errata and Changes in Documentation for JUNOS Release 9.6 for SRX Series Services Gateways on page 143

Issues in JUNOS Release 9.6 for SRX Series Services Gateways

- Outstanding Issues In JUNOS Release 9.6 for SRX Series Services Gateways on page 130
- Resolved Issues in JUNOS Release 9.6 for SRX Series Services Gateways on page 142

Outstanding Issues In JUNOS Release 9.6 for SRX Series Services Gateways

The following problems currently exist in SRX Series Services Gateways. The identifier following the description is the tracking number in the Juniper Networks bug database.

Application Layer Gateways (ALGs)

- On SRX5600 devices, if you run the `show security alg sip counters` command while doing a bulk call generation, it might bring down the SPU with a flowd core file error. [PR/292956]
- On SRX210 devices, the SCCP call cannot be set up after disabling and enabling the SCCP ALG. The call does not go through. [PR/409586]

Authentication

- After the user is authenticated, if 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 Cluster

- Configuring an SRX Series device with the **set system process jsrp-service disable** command only on a primary node of the cluster causes the cluster to go into an incorrect state. [PR/292411]
- The SRX Series device will crash if you use the **set system processes chassis-control disable** command for 4 to 5 minutes and then enable it. Do not use this command on an SRX Series device in a chassis cluster. [PR/296022]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, 8-queue configurations are not reflected on the chassis cluster interface. [PR/389451]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the **iflset** functionality is not supported for aggregated interfaces like **reth**. [PR/391377]
- On SRX210 devices in a chassis cluster, when you upgrade the nodes, sometimes the forwarding daemon might crash and get restarted. [PR/396728]
- On an SRX210 device in a chassis cluster, when you upgrade to the latest software image, the interface links do not come up and are not seen in the Packet Forwarding Engine. As a workaround, you can reboot the device to bring up the interface. [PR/399564]
- On SRX210 devices in a chassis cluster, sometimes the **reth** interface MAC address might not make it to the switch filter table. This results in the dropping of traffic sent to the **reth** interface. As a workaround, restart the Packet Forwarding Engine. [PR/401139]
- On an SRX210 device in a chassis cluster, the fabric monitoring option is enabled by default. This can cause one of the nodes to move to a disabled state. You can disable fabric monitoring by using the following CLI command:


```
set chassis cluster fabric-monitoring disable
```

[PR/404866]
- On the SRX210 Low Memory device in a chassis cluster, the firewall filter does not work on the **reth** interfaces. [PR/407336]
- On SRX210 devices in a chassis cluster, the restart forwarding method is not recommended because when the control link goes through forwarding, the restart forwarding process causes disruption in the control traffic. [PR/408436]
- On an SRX210 device in a chassis cluster, there might be a loss of about 5 packets with 20 Mbps of UDP traffic on an RG0 failover. [PR/413642]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, no trap is generated for redundancy group 0 failover. You can check on the redundancy group 0 state

only when you log in to the device. The nonavailability of this information is caused by a failure of the SNMP walk on the backup/secondary node. As a workaround, use a master-only IP address across the cluster so that you can query a single IP address and that IP address will always be the master for redundancy group 0. [PR/413719]

- On an SRX210 device with an FTP session ramp-up rate of 70, either of the following might disable the secondary node:
 - Back-to-back redundancy group 0 failover
 - Back-to-back primary node reboot
 [PR/414663]
- If an SRX210 device receives more traffic than it can handle, node1 either disappears or gets disabled. [PR/416087]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices in an active/active chassis cluster, when the fabric link fails and then recovers, services with a short time-to-live (such as ALG FTP) stop working. [PR/419095]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices in a chassis cluster, ECMP is not supported for **reth** interfaces. Ping does not go through after configuring load-balance per packet and next hop on **reth**. [PR/423953]
- On SRX3400 and SRX3600 devices in a chassis cluster, ESP authentication errors occur while traffic is sent through 4000 site-to-site IPsec tunnels. [PR/426073]
- On SRX650 devices, doing a redundancy group 0 failover with 1000 logical interfaces on the **reth** interface causes replication errors. As a result, the **ksyncd** process generates a core file. [PR/428636]
- On SRX5800 devices, SNMP traps might not be generated for the ineligible-primary state. [PR/434144]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices in chassis cluster active/active mode, the J-Flow samplings do not happen and the records are not exported to the cflowd server. [PR/436739]
- On SRX5600 and SRX5800 devices, during datapath debugging on a chassis cluster in active/active mode, the IOC EZchip egress trace messages are not traced. [PR/440019]
- On SRX240 Low Memory and High Memory devices, binding the same IKE policy to a dynamic gateway and a site-to-site gateway is not allowed. [PR/440833]
- On SRX650 devices, you will notice this warning messages on the new primary node after a reboot or a RG0 failover:

```
WARNING: cli has been replaced by an updated version:
CLI release 9.6B1.5 built by builder on 2009-04-29 08:24:20 UTC
Restart cli using the new version ? [yes,no] (yes) yes
```

[PR/444470]

- On SRX650 devices in active/active mode, FTP fail transfer might fail after you reboot the active redundancy group node. [PR/454503]
- On SRX240 devices, the cluster might get destabilized when the file system is full and logging is configured on JSRPD and chassisd. The log file size for the

various modules should be appropriately set to prevent the file system from getting full. [PR/454926]

Class of Service (CoS)

- On SRX Series devices, class-of-service-based forwarding (CBF) is not working. [PR/304830]
- On SRX5600 devices, CoS is not supported in transparent mode. [PR/424286]

Flow and Processing

- 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] [PR/397300]
- Configuring the flow filter on SRX Series devices 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]
- On SRX210 and SRX240 devices, broadcast TFTP is not supported when `flow` is enabled on the device. [PR/391399]
- On SRX210, SRX240, and SRX650 devices, after the device fragments packets, the FTP over a GRE link might not perform properly due to packet serialization. [PR/412055]
- On SRX240 devices, traffic flooding happens when multiple multicast IP group addresses are mapped to the same multicast MAC address because multicast switching is based on the Layer 2 address. [PR/418519]
- On SRX650 devices, the input DA errors are not updated when packets are dropped due to MAC filtering on the following:
 - SRX240
 - SRX210
 - 16-port and 24-port GPIMs
 - SRX650 front-end port

This is due to MAC filtering implemented in hardware.

[PR/423777]

- On SRX5600 and SRX5800 devices, the network processing bundle configuration CLI does not check if PICs in the bundle are valid. [PR/429780]
- On SRX650 devices packet loss is observed when the device interoperates with an SSG20 device with AMI line-encoding. [PR/430475]

- On an SRX210 on-board Ethernet port, an IPv6 multicast packet received gets duplicated at the ingress. This happens only for IPv6 multicast traffic in ingress. [PR/432834]
- On SRX3400 and SRX3600 devices, the ramp rate of session creation is slow at times for fragmented UDP traffic. [PR/434508]
- On SRX5800 devices, when there are nonexistent PICs in the network processing bundle, the traffic is sent out to the PICs and is lost. [PR/434976]
- On SRX5800 devices, network processing bundling is not supported in Layer 2 transparent mode. [PR/436863]
- The SRX5600 and SRX5800 devices create more than expected flow sessions with NAT traffic. [PR/437481]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, there is missing information in the `jnxJsFwAuthMultipleFailure` trap message. The trap message is required to contain the Username, IP address, Application, and Trap name whereas, the Username is missing. [PR/439314]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, before initial primary reboot during Low-Impact Cluster Upgrade (LICU), ISSU abort messages are seen even though ISSU is successful. [PR/440545]
- On SRX5800 devices, for any network processing bundle configuration change to take effect, a reboot is needed. Currently there is no message displayed after a bundle configuration change. [PR/441546]
- On SRX5800 devices, the IOC hot swap is not supported with network processing bundling. If an IOC that has a network processing bundling configured gets unplugged, all traffic to that network processor bundle will be lost. [PR/441961]
- On SRX5800 devices with interfaces in a network processing bundle, the ICMP flood or UDP flood cannot be detected at the threshold rate. However, it can be detected at a higher rate when per network processor rate reaches the threshold. [PR/442376]
- On SRX3400 devices in combo mode with 2 SPCs and 1 NPC, not all sessions are created under stress test. [PR/450482]
- On SRX5600 and SRX5800 devices with datapath debugging enabled, multicast packets are not traced at IOC egress chip. [PR/455608]
- On SRX5600 and SRX5800 devices, system log messages are not generated when CPU utilization returns to normal. [PR/456304]

Hardware

- On SRX210 devices, the MTU size is limited to 1518 bytes for the 1-port SFP Mini-PIM. [PR/296498]
- On SRX240 and SRX650 devices and 16-port or 24-port GPIMs, the 1G half-duplex mode of operation is not supported in the autonegotiation mode. [PR/424008]
- On SRX240 devices, the Mini-PIM LEDs glow red for a short duration (1 second) when the device is powered on. [PR/429942]
- On SRX240 devices, the file installation fails on the right USB slot when both of the USB slots have USB keys attached. [PR/437563]

- On SRX240 devices, the combinations of Mini-PIMs cause SFP-Copper links to go down in some instances during bootup, restarting fwdd, and restarting chassisd. As a workaround, reboot the device and the link will be up. [PR/437788]
- On SRX240 devices, using plug and play or booting from Sandisk USB keys can cause a kernel crash. [PR/475315]

Interfaces and Routing

- On SRX3400 devices, the IPv6 transit counters on the **reth** interface show invalid value statistics. [PR/391407]
- On SRX5600 and SRX5800 devices, ping to far-end **reth** interfaces does not work for different routing instances. [PR/408500]
- On SRX240 devices, drops in out-of-profile LLQ packets might be seen in the presence of data traffic, even when the combined (data + LLQ) traffic does not oversubscribe the multilink bundle. [PR/417474]
- On SRX240 and SRX650 devices, when you are configuring the link options on an interface, only the following scenarios are supported:
 - Autonegotiation is enabled on both sides.
 - Autonegotiation is disabled on both sides (forced speed), and both sides are set to the same speed and duplex.

If one side is set to autonegotiation mode and the other side is set to forced speed, the behavior is indeterminate and not supported. [PR/423632]
- On SRX Series devices, the RPM operation will not work for the probe-type **tcp-ping** when the probe is configured with the option **destination-interface**. [PR/424925]
- On SRX650 devices, the following loopback features are not implemented for T1/E1 GPIMs:
 - Line
 - FDL payload
 - Inband line
 - Inband payload

[PR/425040]
- On SRX650 devices, configuring dual and quad T1/E1 framing at the chassis level has no effect. [PR/432071]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, interface statistics on the **st0** interface are not accurate. As a workaround, use the statistics on the security association (SA) to determine input and output bytes and packets. [PR/436857]
- On SRX240 devices, the serial interface maximum speed in extensive output is displayed as 16384 Kbps instead of 8.0 Mbps. [PR/437530]
- On SRX Series devices, incorrect Layer 2 circuit replication on the backup Routing Engine might occur when you:

- Configure nonstop active routing (NSR) and Layer 2 circuit standby simultaneously and commit them
- Delete the NSR configuration and then add the configuration back when both the NSR and Layer 2 circuits are up

As a workaround:

1. Configure the Layer 2 circuit for non-standby.
2. Change the configuration to standby.
3. Add the NSR configuration.

[PR/440743]

- On SRX210 Low Memory devices, the E1 interface will flap and traffic will not pass through the interface if you restart forwarding while traffic is passing through the interface. [PR/441312]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, when you configure the SAP listen option using the `protocol sap listen` command in the CLI, listening fails in both sparse and sparse-dense modes. [PR/441833]
- On SRX100 and SRX210 devices, out-of-band dial-in access using a serial modem does not work. [PR/458114]

Intrusion Detection and Prevention (IDP)

- 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]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, when the device is processing heavy traffic, the `show security idp status operational` command might fail. As a result, IDP flow, session, and packet statistics do not match firewall statistics. [PR/389501] [PR/388048]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, when performing SSL inspection, the HTTPS sessions with higher data transaction sizes fail due to heavy CPU usage. As a result, new connections might fail. [PR/390308]
- SRX100, SRX210, SRX240, and SRX650 devices support only one IDP policy at any given time. When you make changes to the IDP policy and commit, the current policy is completely removed before the new policy becomes effective. During the update, IDP will not inspect the traffic that is passing through the device for attacks. As a result, there is no IDP policy enforcement. [PR/392421]
- On SRX210, SRX3400, SRX3600, SRX5600, and SRX5800 devices, in J-Web selecting **Configuration > Quick Configuration > Security Policies > IDP Policies > Security Package Update > Help** brings up the IDP policy Help page instead of the Signature update Help page. To access the corresponding Help page, select: **Configuration > Quick Configuration > IDP Policies > Signature/Policies Update** and then click Help. [PR/409127]
- On SRX210 devices, during attack detection, multiple attacks get detected. This happens when the IDP policy contains rules that have the match criteria for the

same attacks. Error/warning messages do not appear during policy compilation. [PR/414416]

- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, if you want to change to dedicated mode, the configuration of the **security forwarding-process application-services maximize-idp-sessions** command should be done right before rebooting the device. This should be done to avoid recompiling IDP policies during every commit. [PR/426575]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, IDP is not officially supported in an active/active chassis cluster configuration. The user must disable the IDP configuration when the devices are configured in an active/active chassis cluster. [PR/432252]
- When you configure IDP to run in decoupled mode using the **set security forwarding-process application-services maximize-idp-sessions** command, network address translation (NAT) information will not be shown in the event log. [PR/445908]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, if you configure a policy containing more than 70 rules, with each rule containing the predefined attack groups (Critical, Major, and Minor), the memory constraint of the Routing Engine (500 MB) is reached. [PR/449731]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, during stress conditions one of the SPUs might return an error to the Routing Engine IDP CLI. Currently there is a defect in the Routing Engine IDP CLI that sometimes causes the attack table counters to be shown as a very large (incorrect) value. As a workaround, retry the CLI command to display the correct value. [PR/455544]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices in maximize-idp-sessions mode, there is an IPC channel between two data plane processes. The channel is responsible for transferring the "close session" message (and some other messages) from the FW process to the IDP process. Under stress conditions, the channel is full and extra messages might get lost. This causes IDP sessions in the IDP process to hang, but they will time out eventually. [PR/458900]

J-Flow

- SRX3400, SRX3600, SRX5600, and SRX5800 devices support 4-byte autonomous system (AS) for BGP configuration. However, the J-Flow template versions 5 and 8 do not support 4-byte AS, because these J-Flow templates have 2 bytes for the SRC/DST AS field. [PR/416497]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, J-Flow sampling on the virtual router interface does not show the values of autonomous system (AS) and mask length values. The AS and mask length values of **cflowd** packets show 0 while sampling the packet on the virtual router interface. [PR/419563]
- On SRX Series devices, J-Flow Multicast traffic is not sampled in the output direction. [PR/447357]

J-Web

- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the LEDs on the Routing Engine and PICs are not shown as green when they are up and online on the J-Web Chassis View. [PR/297693]
- On SRX Series devices, when the user adds LACP interface details, a pop-up window appears in which there are two buttons to move the interface left and right. The LACP page currently does not have images incorporated with these two buttons. [PR/305885]
- On SRX210 devices, there is no maximum length limit when the user commits the hostname in CLI mode; however, only a maximum of 58 characters are displayed in the J-Web **System Identification** panel. [PR/390887]
- On SRX210, SRX240, and SRX650 devices, the complete contents of the ToolTips are not displayed in the J-Web Chassis View. As a workaround, drag the Chassis View image down to see the complete ToolTip. [PR/396016]
- On SRX100, SRX210, SRX240, and SRX650 devices, the LED status in the Chassis View is not in sync with the LED status on the device. [PR/397392]
- On SRX Series devices, when you right-click **Configure Interface** on an interface in the J-Web Chassis View, the **Configure > Interfaces** page for all interfaces is displayed instead of the configuration page for the selected interface. [PR/405392]
- On SRX210, SRX3400, SRX3600, SRX5600, and SRX5800 devices, selecting **Configure > Security > Policy > IDP Policies > Security Package Update > Help** in the J-Web user interface brings up the IDP policy Help page instead of the Signature update Help page. To access the corresponding Help page, select **Configure > IDP > Signature Update** and then click **Help**. [PR/409127]
- On SRX Series devices, the CLI Terminal feature is not working in J-Web over IPv6. [PR/409939]
- On all SRX Series devices, IDP Custom Attacks and Dynamic Attack groups cannot be configured using J-Web. [PR/416885]
- On SRX Series devices, it might take extra time to load the J-Web pages when you click **Add** or **Edit** in the STP, GVRP and IGMP-Snooping configuration pages. [PR/422523]
- On SRX210 and SRX240 devices, when J-Web users select the tabs on the bottom-left menu, the corresponding screen is not displayed fully, so users must scroll the page to see all of the content. This issue occurs when the computer is set to a low resolution. As a workaround, set the computer resolution to 1280 x 1024. [PR/423555]
- On SRX Series devices, users cannot differentiate between Active and Inactive configurations on the System Identity, Management Access, User Management, and Date & Time pages. [PR/433353]
- On SRX210 devices, in the J-Web Chassis View, right-clicking any of the ports and clicking **Configure Port** takes the user to the Link aggregation page. [PR/433623]

- On SRX100 devices, in J-Web users can configure the scheduler without entering any stop date. The device submits the scheduler successfully, but the submitted value is not displayed on screen or saved in the device. [PR/439636]
- On SRX100, SRX210, SRX240, and SRX650 devices, in J-Web the associated **dscp** and **dscpv6** classifiers for a logical interface might not be mapped properly when the user edits the classifiers of a logical interface. This can affect the Delete functionality as well. [PR/455670]
- On SRX100, SRX210, SRX240, and SRX650 devices, when J-Web is used to configure a VLAN, the option to add an IPv6 address appears. Only IPv4 addresses are supported. [PR/459530]
- On SRX Series devices in J-Web the left-side menu items and page content might disappear when Troubleshoot is clicked twice. As a workaround, click the Configure or Monitor menu to get back the relevant content. [PR/459936]
- On SRX100, SRX210, SRX240, SRX650 devices, in J-Web, the options Input Filter and Output Filter are displayed in the VLAN configuration page. This feature is not supported, and the user cannot obtain or configure any value under these filter options. [PR/460244]

Management and Administration

- On SRX3400 and SRX3600 devices, the minor alarm is not triggered when the central point or SPU session table is full. [PR/405990]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the queue statistics are not correct after deletion and re-creation of a logical interface (IFL) or creation of a new IFL. IFL statistics are not cleared for 15 minutes after chassis-control is restarted. [PR/417947]
- On SRX5600 devices, when the system is in an unstable state (for example SPU reboot), NFS might generate **residual.nfs** files under the **/var/tmp** directory, which can occupy the disk space for a very long time. As a workaround, run the **request sys storage cleanup** command to clean up when the system has low disk space. [PR/420553]
- On SRX5800 devices, when VPN is not in use, the device will not generate the **var/tmp/spu_kmd_init/** file, which is logged by **lked_cfg**. This should not happen because it is not an error condition. As a result disk space might be wasted over time. As a workaround, run the **cp /dev/null /var/tmp/spu_kmd_init** command from the shell to create this file. Also run **request sys storage cleanup** to clean up when the system has low disk space. [PR/425380]
- On SRX650 devices, the kernel crashes when the link goes down during TFTP installation of the **srxsme** image. [PR/425419]
- On SRX650 devices, continuous messages are displayed from **syslogd** when ports are in switching mode. [PR/426815]
- On SRX240 devices, if a timeout occurs during the TFTP installation, booting the existing kernel using the boot command might crash the kernel. As a workaround, use the reboot command from the loader prompt. [PR/431955]
- On SRX240 devices, when you configure the system log hostname as 1 or 2, the device goes to the shell prompt. [PR/435570]

- On SRX240 devices, the **Scheduler Oinker** messages are seen on the console at various instances with various Mini-PIM combinations. These messages are seen during bootup, restarting fwdd, restarting chassisd, and configuration commits. [PR/437553]
- On SRX5600 and SRX5800 devices, during datapath debugging, the IPsec packets are not traced at the IOC EZchip egress event. [PR/441663]
- On an SRX Series device with **session-init** and **session-close** enabled, do not clear sessions manually when too many sessions are in status "used". [PR/445730]
- On SRX5600 and SRX5800 devices, datapath debug trace messages are dropped at above 1000 packets per second (pps). [PR/446098]
- On SRX5600 and SRX5800 devices, when performing datapath debugging on a chassis cluster in active/active mode, IOC ingress/egress trace messages always show a CID value as 01. The value should be 01 for ingress and 02 for egress if the traffic is entering at node0. [PR/451308]

Power over Ethernet (PoE)

- On SRX240 and SRX210 devices, the output of the PoE operational commands takes roughly 20 seconds to reflect a new configuration or a change in status of the ports. [PR/419920]
- On SRX210 and SRX240 devices, the **deactivate poe interface all** command does not deactivate the PoE ports. Instead, the PoE feature can be turned off by using the **disable** configuration option. Otherwise, the device must be rebooted for the deactivate setting to take effect. [PR/426772]
- On SRX210 and SRX240 devices, the output for the **show poe telemetries** command shows the telemetry data in chronological order. This should be changed to reverse-chronological (most recent data first). [PR/429033]
- On SRX210 and SRX240 devices, the class-4 powered device does not get powered on when PoE is configured to operate in class management mode. [PR/437406]
- The SRX210 and SRX240 devices, the powered device takes more time than what is specified by the standards to power off when operating under overload conditions. [PR/437416]
- On SRX240 and SRX210 devices, the last powered device will not power on if the allocated power becomes equal to the power limit on the device. Power allocated must always be less than the power limit. For example, on the SRX240 device, the powered devices cannot be configured such that allocated power becomes 150 W, even though it is possible to allocate the power up to 149.8 W. [PR/437792]
- On SRX210 and SRX240 devices, reset of the POE controller fails when the **restart chassis-control** command is issued and also after system reboot. PoE functionality is not negatively impacted by this failure. [PR/441798]

Security

- The SRX Series devices do not support egress filter-based forwarding (FBF). [PR/396849]

- On SRX210, SRX3400, SRX3600, SRX5600, and SRX5800 devices in a chassis cluster, if the Infranet Controller auth table mapping action is configured as **provision auth table as needed**, UAC terminates the existing sessions after Routing Engine failover. You might have to initiate new sessions. Existing sessions will not get affected after Routing Engine failover if the Infranet Controller auth table mapping action is configured as **always provision auth table**. [PR/416843]

Unified Threat Management (UTM)

- Content filtering provides the ability to block protocol commands. In some cases, blocking these commands interferes with protocol continuity, causing the session to hang. For instance, blocking the **FETCH** command for the IMAP protocol causes the client to hang without receiving any response. [PR/303584]
- The express antivirus initial database download fails due to the slow start of the device interface. To get a proper update, you can either wait until the next auto-update or manually update the database by using the CLI. [PR/388535]
- When the content filtering message type is set to **protocol-only**, customized messages appear in the log file. [PR/403602]
- The express antivirus feature does not send a replacement block message for HTTP upload (POST) transactions if the current antivirus status is **engine-not-ready** and the fallback setting for this state is **block**. An empty file is generated on the HTTP server without any **block** message contained within it. [PR/412632]
- On SRX240 and SRX650 devices, Outlook Express is sending infected mail (with an EICAR test file) to the mail server (directly, not through DUT). Eudora 7 is using the IMAP protocol to download this mail (through DUT). Mail retrieval is slow, and the EICAR test file is not detected. [PR/424797]
- On SRX650 devices operating under stress conditions, the UTM subsystem file partition might fill up faster than UTM can process and clean up existing temporary files. In that case, the user might see error messages. As a workaround, reboot the system [PR/435124]
- On SRX240 devices, FTP download for > 4MB files does not work in a two-device topology. [PR/435366]
- On SRX210, SRX240, SRX650 devices, the Websense server stops taking new connections after HTTP stress. All new sessions get blocked. As a workaround, reboot the Websense server. [PR/435425]
- On SRX240 devices, if the device is under UTM stress traffic for several hours, users might get the following error while using a UTM command:

the utmd subsystem is not responding to management requests.

As a workaround, restart the **utmd** process. [PR/436029]
- The UTM antivirus feature might fail for HTTP scans on files larger than 16 MB. As a workaround, users should set **max-content-size** in the antivirus configuration to 16 MB. Files larger than 16 MB will either be passed or dropped without being scanned, depending on the fallback option that had been set for **max-content-size**. [PR/457472]

Virtual LANs (VLANs)

- On SRX650 devices, when VLAN tagging is configured and traffic is sent, the output of `show interfaces ge-0/0/1 media detail` VLAN tagged frame count is not shown. [PR/397849]
- On SRX240 and SRX650 devices, tagged frames on an access port with the same VLAN tag are not getting dropped. [PR/414856]
- On SRX100 devices, the packets are not being sent out of the physical interface when the VLAN ID associated with the VLAN interface is changed. As a workaround, you need to clear the ARP. [PR/438151]

VPN

- The shared IKE limit for IKE users is not currently enforced. More users than are specified in the shared IKE limit are able to establish IKE/IPsec tunnels. [PR/288551]
- On SRX210 and SRX240 devices, concurrent login to the device from a different management systems (for example, laptop or computers) are not supported. The first user session will get disconnected when a second user session is started from a different management system. Also, the status in the first user system is displayed incorrectly as “Connected”. [PR/434447]
- In SRX Series devices, the site-to-site policy-based VPNs in a three or more zone scenario will not work if the policies match the address 'any', instead of specific addresses, and all cross-zone traffic policies are pointing to the single site-to-site VPN tunnel. As a workaround, configure address books in different zones to match the source and destination, and use the address book name in the policy to match the source and destination. [PR/441967]

Resolved Issues in JUNOS Release 9.6 for SRX Series Services Gateways

The following issues have been resolved since JUNOS Release 9.6 R2. The identifier following the description is the tracking number in our bug database.

Chassis Cluster

- On SRX5600 and SRX5800 devices in a chassis cluster, whenever the `reth` interface with static MAC addresses was configured, the ping operation failed from the directly connected device to the chassis cluster. [PR/455051: This issue has been resolved.]
- On SRX3400, SRX3600, SRX5600, and SRX5800 devices, the firewall filter stopped working after you rebooted the primary node in a cluster multiple times. [PR/456116: This issue has been resolved.]

Flow and Processing

- On SRX650 devices in which traffic streams had similar source MAC addresses, the hashkey ended up being the same. Because of this, all the streams got mapped to the same single uplink. As a result, the uplink to the CPU was exhausted and the system was limited to 2.5 GB throughput traffic. [PR/428526: This issue has been resolved.]
- On an SRX5800 device with a 1-Gbps IOC, when more than 10 ports per port module were used, intermittent packet loss occurred because of oversubscription. [PR/433209: This issue has been resolved.]

Intrusion Detection and Prevention (IDP)

- An IDP policy failed to load if the IDP policy name had spaces or tabs. [PR/493007: This issue has been resolved.]
- When IDP inspection was enabled, traffic outage was observed intermittently because of a memory corruption in Juniper Pattern Matching Engine (JPME) queues. [PR/496052: This issue has been resolved.]

J-Web

- On SRX Series devices, on the J-Web spanning-tree configuration page, the Edit interface/msti window did not save the data before committing the configuration. [PR/433506: This issue has been resolved.]

Virtual LANs (VLANs)

- On SRX100 and SRX240 devices, default VLAN was not added to the Switch Trunk with "VLAN member all" configuration after reboot. The default VLAN was only used for access ports, to quickly set up the device to switch traffic. [PR/450869: This issue has been resolved.]

Related Topics

- New Features in JUNOS Release 9.6 for SRX Series Services Gateways on page 110
- Known Limitations in JUNOS Release 9.6 for SRX Series Services Gateways on page 125
- Errata and Changes in Documentation for JUNOS Release 9.6 for SRX Series Services Gateways on page 143

Errata and Changes in Documentation for JUNOS Release 9.6 for SRX Series Services Gateways

This section lists outstanding issues with the documentation.

Application Layer Gateways (ALGs)

- ALG configuration examples in the *JUNOS Software Security Configuration Guide* incorrectly show policy-based NAT configurations. NAT configurations are now rule-based.

Attack Detection and Prevention

- The default parameters documented in the firewall/NAT screen configuration options table in the *JUNOS Software Security Configuration Guide* and the J-Web online Help do not match the default parameters in the CLI. The correct default parameters are:

```
tcp {
  syn-flood {
    alarm-threshold 1024;
    attack-threshold 200;
    source-threshold 1024;
    destination-threshold 2048;
    timeout 20;
  }
}
[edit security screen ids-option untrust-screen]
```

Chassis Cluster

- The current chassis cluster documentation does not include the following information:

When performing back-to-back redundancy group 0 failovers, leave at least 5 minutes between failovers.
- Some information about ISSU initiation and troubleshooting was left out of the *JUNOS Software Security Configuration Guide*. There is also one inaccuracy in the ISSU information. The first three items below are information missing from the guide. The fourth item explains the mistake.

- **Reboot required**— In Step 2 of the ISSU initiation procedure described in the *JUNOS Software Security Configuration Guide*, the command is incomplete. The end of the command should include the **reboot** keyword. If **reboot** is not included in the command, you will need to manually reboot each device as the ISSU completes updating the software image. The “request system software in-service-upgrade” section of the *JUNOS Software CLI Reference* also fails to note the **reboot** requirement.

The following example shows the correct ISSU initiation command sequence.

```
user@host> request system software in-service-upgrade image_name reboot
```

- **Automatic fallback**— If you want redundancy groups to automatically return to node 0 as the primary after the ISSU is complete, you must set the redundancy group priority such that node 0 is primary and enable the preempt option. Note that this method will work for all redundancy groups except redundancy group 0. You must manually fail over redundancy group

0. To set the redundancy group priority and enable the preempt option, see “Configuring Redundancy Groups” in the *JUNOS Software Security Configuration Guide*.

To manually fail over a redundancy group, see “Initiating a Manual Redundancy Group Failover” in the *JUNOS Software Security Configuration Guide*.

- **Roll back one device**—If the ISSU fails to complete and only one device in the cluster has been upgraded, you can roll back to the previous configuration on that device alone by using the following commands on the upgraded device:

1. `request chassis cluster in-service-upgrade abort`
2. `request system software rollback`
3. `request system reboot`

There is currently no method to roll back both devices without a service disruption.

- **GRES error in ISSU descriptions**—There is a mistake in item 3 of the “Before You Begin” list in the “Low-Impact ISSU Chassis Cluster Upgrades” section of the *JUNOS Software Security Configuration Guide*. Item 3 states, “We also recommend that graceful Routing Engine switchover (GRES) be enabled prior to starting an ISSU.” The item should read, “We also recommend that routing protocols graceful restart be enabled prior to starting an ISSU.”

The same error is also present in the `request system software in-service-upgrade` command description of the *JUNOS Software CLI Reference*.

CLI Reference

The “Services Configuration Statement Hierarchy” section in the *JUNOS Software CLI Reference* refers to the *JUNOS Services Interfaces Configuration Guide*, which has the following error in the sections “Data Size” and “Configuring the Probe”:

- The minimum data size required by the UDP timestamp probe is identified as 44 bytes. This is incorrect: the minimum data size required by the UDP timestamp probe is 52 bytes.

CompactFlash Card Support

- The *JUNOS Software Administration Guide* incorrectly states that JUNOS supports a 256-MB CompactFlash card size. JUNOS supports only 512-MB and 1024-MB CompactFlash card sizes.

DLSw

- The *JUNOS Software Interfaces and Routing Configuration Guide* incorrectly states that the data link switching (DLSw) protocol is supported in this release. DLSw support ended in JUNOS Release 9.3.

Flow

- The *Junos OS CLI Reference* and *Junos OS Security Configuration Guide* state that the following aggressive aging statements are supported on SRX Series devices when in fact they are not supported on SRX3400, SRX3600, SRX5600, and SRX5800 devices:
 - [edit security flow aging early-ageout]
 - [edit security flow aging high-watermark]
 - [edit security flow aging low-watermark]
- The “Understanding Selective Stateless Packet-Based Services” section in the *JUNOS Software Administration Guide* states: “The following security features are not supported with selective stateless packet-based services—stateful firewall NAT, IPsec VPN, DOS screens, J-flow traffic analysis, WXC integrated security module, security policies, zones, attack detection and prevention, PKI, ALGs, and chassis cluster.” This statement is not correct. With selective packet-mode, traffic that is sent through flow is able to use all of those services, even in a single VR scenario.

Incorrect Administration Features Support Information in Documentation

The *JUNOS Software Administration Guide* lists the following incorrect support information in Table 35: Administrator Authentication in “Administration Features on SRX3400, SRX3600, SRX5600, and SRX5800 Services Gateways”:

Table 35: Administrator Authentication

Feature	More information
Local authentication (SRX5600 and SRX5800 only)	User Authentication Overview

Local authentication is also supported on SRX3400 and SRX3600 devices.

Incorrect Security Features Support Information in Documentation

The *JUNOS Software Security Configuration Guide* lists the following incorrect support information in Table 39: IPsec in “Security Features on SRX3400, SRX3600, SRX5600, and SRX5800 Services Gateways”:

Table 39: IPsec

Feature	More information
XAuth extended authentication for remote access connections (SRX5600 and SRX5800 only)	User Authentication Overview
VPN monitoring (SRX5600 and SRX5800 only)	Configuring VPN Global Settings (Standard VPNs)

Both XAuth extended authentication for remote access connections and VPN monitoring are also supported on SRX3400 and SRX3600 devices.

Installing Software Packages

- The current SRX210 documentation does not include the following information:
On SRX210 devices, the */var* hierarchy is hosted in a separate partition (instead of the *root* partition). If JUNOS software installation fails as a result of insufficient space:
 1. Use the `request system storage cleanup` command to delete temporary files.
 2. Delete any user-created files in both the *root* partition and under the */var* hierarchy.
- The “Installing Software using the TFTPBOOT Method on the SRX100, SRX210, and SRX650 Services Gateway” section in the *JUNOS Software Administration Guide* contains the following inaccuracies:

- The documentation incorrectly implies that the TFTPBOOT method requires a separate secondary device to retrieve software from the TFTP server.
- The documentation should indicate that the TFTPBOOT method does not work reliably over slow speeds or large latency networks.
- The documentation indicates that before starting the installation, you only need to configure the gateway IP, device IP address, and device IP netmask manually in some cases, when actually you need to configure them manually in all cases.
- The documentation should indicate that on the SRX100, SRX210, and SRX240 devices, only the ge-0/0/0 port supports TFTP in uboot and on the SRX650 device, all front-end ports support TFTP in uboot.
- Step 2 of the “Installing JUNOS Software Using TFTPBOOT” instructions should mention that the URL path is relative to the TFTP server’s TFTP root directory. The instructions should also mention that you should store the JUNOS image file in the TFTP server’s TFTP root directory.
- The documentation should indicate that the TFTPBOOT method installs software on the internal flash on SRX100, SRX210, and SRX240 devices, whereas on SRX650 devices, the TFTP method can install software on the internal or external CompactFlash card.
- The *JUNOS Software Administration Guide* is missing the following information about installing software using USB on SRX100, SRX210, SRX240, and SRX650 devices:

You can install or recover the JUNOS software using USB on SRX100, SRX210, SRX240, and SRX650 devices. During the installation process, the installation package from the USB is installed on the specified boot media.

Before you begin the installation, ensure the following prerequisites are met:

- U-boot and Loader are up and running on the device.
- USB is available with the JUNOS package to be installed on the device.

To install the software image on the specified boot media:

1. Go to the Loader prompt. For more information on accessing the Loader prompt, see “Accessing the Loader Prompt” on page 260 of the *JUNOS Software Administration Guide*.
2. Enter the following command at the Loader prompt:

```
Loader > install URL
```

Where URL is *file:///package*

Example:

```
Loader > install file:///junos-srxsme-9.4-200811.0-domestic.tgz
```

When you are done, the file reads the package from the USB and installs the software package. After the software installation is complete, the device boots from the specified boot media.



NOTE: USB to USB installation is not supported. Also, on SRX100, SRX210, and SRX240 devices, the software image will always be installed on NAND flash, but on SRX650 devices, the software image can be installed either on the internal or external CompactFlash card based on the boot media specified.

Intrusion Detection and Prevention (IDP)

- In the *JUNOS Software Security Configuration Guide*, the following information in the "Verifying the Policy Compilation and Load Status" section is incorrect:
 - The text does not indicate that the log file must be created first.
 - The path for the log file is incorrect.

Note the following correct information:

- Create the log file first by entering `set security idp traceoptions file idpd`. You can then set flags by entering `set security idp traceoptions flag all`.
- The correct path for the idpd log file is `/var/log`, not `/var/db`
- The "Configuring SSL Inspection" section of the *JUNOS Software Security Configuration Guide* incorrectly states that SSL inspection is disabled by default and is enabled if any configurations are detected. This information is obsolete as of JUNOS Release 10.2. Although previously the SSL decoder was disabled by default in the detector, as of JUNOS Release 10.2 the SSL decoder is enabled by default. The updated information is as follows: The SSL decoder is enabled by default. To manually enable it, use the following CLI command:

```
user@host>set security idp sensor-configuration detector protocol-name
SSL tunable-name sc_ssl_flags tuneable-value 1
```

- The IDP rule notification options listed in the *JUNOS Software Security Configuration Guide* incorrectly include the **Send Emails** and **Run Scripts** options, which are not supported in the JUNOS 9.6 Release.

J-Web

- **J-Web security package update Help page**—The J-Web Security Package Update Help page does not contain information about download status.
- **J-Web pages for stateless firewall filters**—There is no documentation describing the J-Web pages for stateless firewall filters. To find these pages in J-Web, go to **Configure > Security > Firewall Filters**, then select **IPv4 Firewall Filters** or **IPv6 Firewall Filters**. After configuring filters, select **Assign to Interfaces** to assign your configured filters to interfaces.

Network Management

The *Junos Network Management Configuration Guide* does not include the following information:

- **MIB Objects for the SRX100 Services Gateway**—The Chassis MIB objects for the SRX100 Services Gateway include:

```

jnxProductLineSRX100    OBJECT IDENTIFIER ::= { jnxProductLine 41 }
jnxProductNameSRX100    OBJECT IDENTIFIER ::= { jnxProductName 41 }
jnxChassisSRX100        OBJECT IDENTIFIER ::= { jnxChassis      41 }

jnxSlotSRX100           OBJECT IDENTIFIER ::= { jnxSlot        41 }
jnxSRX100SlotFPC        OBJECT IDENTIFIER ::= { jnxSlotSRX100   1 }
jnxSRX100SlotRE         OBJECT IDENTIFIER ::= { jnxSlotSRX100   2 }
jnxSRX100SlotPower      OBJECT IDENTIFIER ::= { jnxSlotSRX100   3 }
jnxSRX100SlotFan        OBJECT IDENTIFIER ::= { jnxSlotSRX100   4 }

jnxMediaCardSpaceSRX100 OBJECT IDENTIFIER ::= { jnxMediaCardSpace
41 }
jnxSRX100MediaCardSpacePIC OBJECT IDENTIFIER ::= {
jnxMediaCardSpaceSRX100 1 }

jnxMidplaneSRX100       OBJECT IDENTIFIER ::= { jnxBackplane 41 }

jnxModuleSRX100         OBJECT IDENTIFIER ::= { jnxModule     41 }
jnxSRX100FPC            OBJECT IDENTIFIER ::= { jnxModuleSRX100 1 }
jnxSRX100RE             OBJECT IDENTIFIER ::= { jnxModuleSRX100 2 }
jnxSRX100Power          OBJECT IDENTIFIER ::= { jnxModuleSRX100 3 }
jnxSRX100Fan            OBJECT IDENTIFIER ::= { jnxModuleSRX100 4 }

```

- **MIB Objects for the SRX210 Services Gateway**—The Chassis MIB objects for the SRX210 Services Gateway include:

```

jnxProductLineSRX210    OBJECT IDENTIFIER ::= { jnxProductLine    36
}
jnxProductNameSRX210    OBJECT IDENTIFIER ::= { jnxProductName    36
}
jnxChassisSRX210        OBJECT IDENTIFIER ::= { jnxChassis        36
}

jnxSlotSRX210           OBJECT IDENTIFIER ::= { jnxSlot          36 }
jnxSRX210SlotFPC        OBJECT IDENTIFIER ::= { jnxSlotSRX210     1 }
jnxSRX210SlotRE         OBJECT IDENTIFIER ::= { jnxSlotSRX210     2 }
jnxSRX210SlotPower      OBJECT IDENTIFIER ::= { jnxSlotSRX210     3 }
jnxSRX210SlotFan        OBJECT IDENTIFIER ::= { jnxSlotSRX210     4 }

jnxMediaCardSpaceSRX210 OBJECT IDENTIFIER ::= { jnxMediaCardSpace
36 }
jnxSRX210MediaCardSpacePIC OBJECT IDENTIFIER ::= {
jnxMediaCardSpaceSRX210 1 }

jnxMidplaneSRX210       OBJECT IDENTIFIER ::= { jnxBackplane 36 }

jnxModuleSRX210         OBJECT IDENTIFIER ::= { jnxModule       36 }
jnxSRX210FPC            OBJECT IDENTIFIER ::= { jnxModuleSRX210 1 }
jnxSRX210RE             OBJECT IDENTIFIER ::= { jnxModuleSRX210 2 }
jnxSRX210Power          OBJECT IDENTIFIER ::= { jnxModuleSRX210 3 }
jnxSRX210Fan            OBJECT IDENTIFIER ::= { jnxModuleSRX210 4 }

```

Screens

- The following guide contains incorrect screen configuration instructions:
 - *JUNOS Software Design and Implementation Guide*, “Implementing Firewall Deployments for Branch Offices” chapter

Examples throughout this guide 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 Release 9.6 for SRX Series Services Gateways on page 110
 - Known Limitations in JUNOS Release 9.6 for SRX Series Services Gateways on page 125
 - Issues in JUNOS Release 9.6 for SRX Series Services Gateways on page 130

JUNOS Software Release Notes for Juniper Networks J Series Services Routers

- New Features in JUNOS Release 9.6 for J Series Services Routers on page 152
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series Services Routers on page 156
- Known Limitations in JUNOS Release 9.6 for J Series Services Routers on page 157
- Issues in JUNOS Release 9.6 for J Series Services Routers on page 159
- Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers on page 165
- Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers on page 168
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170

New Features in JUNOS Release 9.6 for J Series Services Routers

The following features have been added to JUNOS Release 9.6. Following the description is the title of the manual or manuals to consult for further information.

- Software Features on page 153

Software Features

Class of Service (CoS)

- **Class-of-service (CoS) components**—CoS components for Asynchronous Transfer Mode (ATM) are provided to control data transfer, especially for time-sensitive voice packets. J Series devices with ADSL Annex A or Annex B PIMs can use an ATM interface to send network traffic through a point-to-point connection to a DSL access multiplexer (DSLAM).

Interfaces and Routing

- **Selective stateless packet-based services**—Allow you to use both flow-based and packet-based forwarding simultaneously on a system. You configure these services by specifying an action modifier in stateless firewall filters and applying the filters to certain interfaces. On these interfaces, traffic that matches the filter criteria will bypass flow-based forwarding. Bypassing flow-based forwarding can be useful for traffic for which you explicitly want to avoid session-scaling constraints.

A defined set of stateless services is available with selective stateless packet-based services.

[Junos OS Administration Guide for Security Devices]

IPsec VPN

- **Hub-and-spoke**—For route-based IPsec VPNs, JUNOS Software supports a topology called hub-and-spoke, which allows two or more remote sites to communicate to each other through a central site using VPN tunnels. To use hub-and-spoke, you create two IPsec VPN tunnels that terminate at a gateway and define a pair of routes so that the device directs traffic exiting from one tunnel to the other tunnel. Only route-based hub-and-spoke is supported.

J-Web

- **J-Web user interface enhancements**—This feature is supported on all J Series devices.

The J-Web user interface has been significantly updated in JUNOS Release 9.6. The menu system has been updated to provide more intuitive navigation to most pages. For example, configuration pages for firewall policies and UTM policies are now grouped together under the Configure > Security > Policy menu. In addition, the look and feel of the monitoring and configuration pages has been updated in JUNOS Release 9.6.

[JUNOS Software Administration Guide, JUNOS Software Interfaces and Routing Configuration Guide, JUNOS Software Security Configuration Guide]

Management and Administration

- **Reverse Telnet**—Reverse Telnet allows you to configure J Series devices to listen on a specific port for Telnet and SSH (secure shell) services. When you connect to that port, the device provides an interface to the auxiliary port on the device. In order to use reverse Telnet, you must have the following: a device with an auxiliary port running the appropriate version of JUNOS Software, and a device with a console port for remote management (if network connectivity fails and you want to use console access). To configure reverse Telnet, use the **reverse telnet port** and **reverse ssh port** statements at the **[set system services]** hierarchy level.

[Junos OS Administration Guide for Security Devices]

- **NAT information in session logs**—As with previous releases, a session can be logged whenever it is created, closed, rejected, or denied. Starting with JUNOS Software Release 9.6, the following additional information is included in session logs for all J Series devices:
 - **nat-source-address**—The translated NAT source address if NAT was applied; otherwise, the source-address.
 - **nat-source-port**—The translated NAT source port if NAT was applied; otherwise, the source-port.
 - **nat-destination-address**—The translated NAT destination address if NAT was applied; otherwise, the destination-address.
 - **nat-destination-port**—The translated NAT destination port (if any); otherwise, the destination-port.
 - **dst-nat-rule-name**—The destination NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if destination address translation takes place, then this field shows the static NAT rule name.
 - **src-nat-rule-name**—The source NAT rule that was applied to the session (if any). If static NAT is also configured and applied to the session and if source address translation takes place, then this field shows the static NAT rule name.



NOTE: Some sessions might have both destination and source NAT applied and the information logged.

- **service-name**—The service (application) through which the packet traversed.
- **session-id-32**—The 32-bit session ID.

Both traditional and structured system log formats are supported, with the exception that structured system logs are only exported to the Routing Engine, they are not exported to an external host.

[Junos OS Security Configuration Guide]

Network Address Translation (NAT)

- **Port randomization for source NAT**—For pool-based source NAT and interface NAT, port numbers are allocated randomly by default. Although randomized port number allocation can provide protection from security threats such as DNS poison attacks, it can also affect performance and memory usage for pool-based source NAT.

You can disable port randomization by using the `port-randomization disable` statement at the `[edit security nat source]` hierarchy level. To re-enable port randomization, use the `port-randomization` statement at the `[edit security nat source]` hierarchy level.

- **Persistent NAT**—Release 9.6 provides configuration of persistent NAT on the J Series Services Routers. Persistent NAT allows applications to use the Session Traversal Utilities for NAT (STUN) protocol for NAT traversal.

The STUN protocol is a client-server protocol that allows a client application behind a NAT device to learn its public IP address and the type of NAT used to allocate the address bindings. The STUN client is commonly used with Session Initiation Protocol (SIP) VoIP applications where IP addresses and port numbers are encoded within the application data. Both the STUN client and STUN server are provided by the application.

Persistent NAT ensures that all requests from the same internal transport address are mapped to the same external transport address by the NAT device closest to the STUN server. The following types of persistent NAT can be configured on the Juniper Networks device:

- **Any remote host**—All requests from a specific internal IP address and port are mapped to the same external IP address and port. Any external host can send a packet to the internal host using the mapped external address when the incoming policy from external to internal is configured.
- **Target host**—All requests from a specific internal IP address and port are mapped to the same external IP address and port. An external host can send a packet to an internal host only if the internal host had previously sent a packet to the external host's IP address.
- **Target host port**—All requests from a specific internal IP address and port are mapped to the same external IP address and port. An external host can send a packet to an internal host only if the internal host had previously sent a packet to the external host's IP address and port.

To configure persistent NAT options, use the `persistent-nat` statement in the `[edit security nat]` hierarchy.

You can configure security policies with two new predefined services, `junos-stun` and `junos-persistent-nat`, to permit or deny persistent NAT traffic.



NOTE: Persistent NAT is sometimes referred to as cone NAT. The term cone NAT has been replaced by persistent NAT by the IETF.

[Junos OS Security Configuration Guide]

Unified Threat Management (UTM)

- **Unified Threat Management J-Web support**—The UTM J-Web Quick Configuration screens have been redesigned to support a new J-Web framework for enhanced usability.

[Junos OS Security Configuration Guide]

- Related Topics**
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series Services Routers on page 156
 - Known Limitations in JUNOS Release 9.6 for J Series Services Routers on page 157
 - Issues in JUNOS Release 9.6 for J Series Services Routers on page 159
 - Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers on page 168
 - Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170

Changes in Default Behavior and Syntax in JUNOS Release 9.6 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 documentation:

Configuration

- J Series devices no longer allow a configuration in which a tunnel's source or destination address falls under the subnet of the same logical interface's address.

Management and Administration

- The following session logging fields have been renamed in JUNOS Release 9.6:

Old Name	New Name
inbound-packets	packets-from-client
inbound-bytes	bytes-from-client
outbound-packets	packets-from-server
outbound-bytes	bytes-from-server

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 `order radius password` or `ldap password`.

Related Topics

- New Features in JUNOS Release 9.6 for J Series Services Routers on page 152
- Known Limitations in JUNOS Release 9.6 for J Series Services Routers on page 157
- Issues in JUNOS Release 9.6 for J Series Services Routers on page 159
- Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers on page 168
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170
- Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers on page 165

Known Limitations in JUNOS Release 9.6 for J Series Services Routers

J Series platforms now support IDP and UTM functionality. Under heavy network traffic in a few areas of functionality, such as NAT and IPsec VPN, performance is still being improved to reach the high levels to which Juniper Networks is consistently committed.

Chassis Cluster

In JUNOS Release 9.6, the following features are not supported when chassis clustering is enabled on the router:

- **Packet-based protocols**—All packet-based protocols, such as 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**
- **J Series chassis cluster**—The maximum number of redundancy groups is equal to the number of redundant Ethernet interfaces configured by the user.

IGMP

- SRX100 devices do not support IGMP snooping.

Interfaces and Routing

- The user can use IPsec only on an interface that resides in the routing instance inet 0. The user will not be able to assign an internal or external interface to the IKE policy if that interface is placed in a routing instance other than inet 0.
- On J Series devices, flow mode does not support asymmetric routing for stateful sessions. As a result of this behavior, trace-route might not work when VRRP is configured across J Series devices.

Intrusion Detection and Prevention (IDP)

- On J Series Services Routers, IP actions do not work when users select a timeout value greater than 65535 in the IDP policy.

J-Web

- Some J-Web pages for new features (for example, the Quick Configuration page for the switching features on J Series Services Routers) display content in one or more modal pop-up windows. In the modal pop-up windows, you can interact only with the content in the window and not with the rest of the J-Web page. As a result, online Help is not available when modal pop-up windows are displayed. You can access the online Help for a feature only by clicking the **Help** button on a J-Web page.

SNMP

- The SNMP NAT-related MIB is not supported in release 9.6.

Unified Threat Management (UTM)

- Unified Threat Management (UTM) requires 1 GB of memory. If your J2320, J2350, or J4350 device has only 512 MB of memory, you must upgrade the memory to 1 GB to run UTM.

Related Topics

- New Features in JUNOS Release 9.6 for J Series Services Routers on page 152
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series Services Routers on page 156
- Issues in JUNOS Release 9.6 for J Series Services Routers on page 159
- Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers on page 168
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170
- Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers on page 165

Issues in JUNOS Release 9.6 for J Series Services Routers

- Outstanding Issues In JUNOS Release 9.6 for J Series Services Routers on page 160
- Resolved Issues in JUNOS Release 9.6 for J Series Services Routers on page 164

Outstanding Issues In JUNOS Release 9.6 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.

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]
- On J4350 Services Routers, 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]
- On J2300, J2320, J2350, J4350, and J6350 Services Routers in a active/active chassis cluster, when the fabric link fails and then recovers, services with a short **time-to-live**, such as ALG FTP, stop working. [PR/419095]
- On J2300, J2320, J2350, J4350, and J6350 Services Routers doing a redundancy group 0 failover with 1000 logical interfaces on the **reth** interface causes replication errors, which makes **ksyncd** generate a core file. [PR/428636]

Class of Service (CoS)

- 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 WXC application acceleration platform might fail over with heavy traffic. [PR/273843]

Enhanced Switching

- 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 and Processing

- 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]
- On J Series Services Routers, NAT traffic that is going to the WXC ISM 200 and returning back in clear (that is, not accelerated by the WXC ISM200) does not work. [PR/438152]

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 device 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 Routing

- 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]
- Asymmetric routing, such as tracing a route to a destination behind J Series routers running JUNOS Software with Virtual Router Redundancy Protocol (VRRP), does not work. [PR/237589]
- On J2320 Services Routers, when you enable the DHCP client, the default route is not added to the route table. [PR/296469]
- On J2320, J2350, J4350, and J6350 Services Routers, broadcast TFTP is not supported when `flow` is enabled on the device. [PR/391399]
- RPM client operation will not work for the `probe-type tcp-ping` when the probe is configured with the option `destination-interface`. [PR/424925]
- In J Series xDSL PIMs, mapping between IP CoS and ATM CoS is not supported. If the user configures IP CoS in conjunction with ATM CoS, the logical interface level shaper matching ATM CoS rate must be configured to avoid congestion drops in SAR.
Example:

```
set interfaces at-5/0/0 unit 0 vci 1.110
set interfaces at-5/0/0 unit 0 shaping cbr 62400 ATM COS
set class-of-service interfaces at-5/0/0 unit 0 scheduler-map sche_map IP COS
set class-of-service interfaces at-5/0/0 unit 0 shaping-rate 62400 ADD IFL SHAPER
```

[PR/430756]
- One member link is going down in a Multilink (ML) bundle during bidirectional traffic with Multilink Frame Relay (MFR). [PR/445679]
- On J Series Services Routers, out-of-band dial-in access using serial modem does not work. [PR/458114]

J-Web

- IDP custom attacks and dynamic attack groups cannot be configured using J-Web. [PR/416885]
- On J Series Services Routers, the `Ajax` calls need to be optimized and should be in synchronization with the existing configuration screens (STP, GVRP, and IGMP snooping). [PR/422523]
- On J2350, J4350, and J6350 Services Routers, when J-Web users select the tabs on the bottom-left menu, the corresponding screen is not displayed fully, so

users must scroll the page to see all content. This issue occurs when the computer is set to a low resolution. As a workaround, set the computer resolution to 1280 x 1024. [PR/423555]

- In the VLAN configuration page, the input filter and output filter combo boxes are displayed, but no filter values will appear, because those options are no longer available. [PR/460244]

Management and Administration

- On a J Series device with **session-init** and **session-close** enabled, do not clear sessions manually when too many sessions are in status "used". [PR/445730]
- On J4350, J6350, J2320 and J2350 devices running in flow-based mode, only the security option **system-services any-service** will allow reverse Telnet and reverse SSH; **system-services all** will not allow reverse Telnet and reverse SSH. [PR/447323]
- Extended Bit Error Rate Test (BERT) takes an additional 3 hours to complete even though a BERT-period of 24 hours is set. [PR/447636]

Unified Access Control (UAC)

- On J Series Services Routers, MAC address-based authentication does not work when the router is configured as a UAC L2 Enforcer. [PR/431595]

Unified Threat Management (UTM)

- On J2320, J2350, J4350, and J6350 Services Routers, Outlook Express is sending infected mail (with an EICAR test file) to a mail server (directly, not through DUT). Eudora 7 is using the IMAP protocol to download this mail (through DUT). Mail retrieval is slow, and the EICAR test file is not detected. [PR/424797]

Virtual LANS (VLANs)

- When using J-Web to configure a VLAN on an SRX or a J Series device, the option to add an IPv6 address appears. Only IPv4 addresses are supported. [PR/459530]

VPN

- Site-to-site policy-based VPNs in a three or more zone scenario will not work if the policies match the address "any", instead of specific addresses, and all cross zone traffic policies are pointing to the single site-to-site VPN tunnel. As a workaround, configure address books in different zones to match the source and destination, and use the address book name in the policy to match the source and destination. [PR/441967]

WXC Integrated Services Module

- When two J Series devices 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 does not support policy-based VPN with WXC Integrated Services Modules (WXC ISM 200s). [PR/281822]

Resolved Issues in JUNOS Release 9.6 for J Series Services Routers

The following issues have been resolved since JUNOS Release 9.6 R2. The identifier following the description is the tracking number in our bug database.

Enhanced Switching

- With a large multiple spanning tree (mstp) configuration, the enhanced-switching-enabled PIC (uPIM) entered an invalid state. In "show chassis fpc pic-status," the uPIM was marked as a "Hardware Error." [PR/462067: This issue has been resolved.]

Flow and Processing

- On J2350, J4350, and J6350 Services Routers, OSPF over GRE over IPsec did not work. [PR/105279: This issue has been resolved.]

J-Web

- On J Series Services Routers, on the spanning-tree configuration page, the **Edit interface/msti** window did not save the data before committing the configuration. [PR/433506: This issue has been resolved.]

Virtual Lans (VLANs)

- On J Series devices, default VLAN was not added to the Switch Trunk with "VLAN member all" configuration after reboot. The default VLAN was only used for access ports, to quickly set up the device to switch traffic. [PR/450869: This issue has been resolved.]

Related Topics

- New Features in JUNOS Release 9.6 for J Series Services Routers on page 152
- Known Limitations in JUNOS Release 9.6 for J Series Services Routers on page 157
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series Services Routers on page 156
- Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers on page 168
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170
- Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers on page 165

Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers

This section lists outstanding issues with the documentation.

Application Layer Gateways (ALGs)

- ALG configuration examples in the *JUNOS Software Security Configuration Guide* incorrectly show policy-based NAT configurations. NAT configurations are now rule-based.

CLI Reference

The “Services Configuration Statement Hierarchy” section in the *JUNOS Software CLI Reference* JUNOS Software CLI Reference refers to the *JUNOS Services Interfaces Configuration Guide*, JUNOS Services Interfaces Configuration Guide, which has the following error in the sections “Data Size” and “Configuring the Probe”:

- The minimum data size required by the UDP timestamp probe is identified as 44 bytes. This is incorrect: the minimum data size required by the UDP timestamp probe is 52 bytes.

DLSw

- The *JUNOS Software Interfaces and Routing Configuration Guide* incorrectly states that the data link switching (DLSw) protocol is supported in this release. DLSw support ended in JUNOS Release 9.3.

Flow

- The “Understanding Selective Stateless Packet-Based Services” section in the *JUNOS Software Administration Guide* states: “The following security features are not supported with selective stateless packet-based services—stateful firewall NAT, IPsec VPN, DOS screens, J-flow traffic analysis, WXC integrated security module, security policies, zones, attack detection and prevention, PKI, ALGs, and chassis cluster.” This statement is not correct. With selective packet-mode, traffic that is sent through flow is able to use all of those services, even in a single VR scenario.

Intrusion Detection and Prevention (IDP)

- In the *JUNOS Software Security Configuration Guide*, the following information in the “Verifying the Policy Compilation and Load Status” section is incorrect:
 - The text does not indicate that the log file must be created first.
 - The path for the log file is incorrect.

Note the following correct information:

- Create the log file first by entering `set security idp traceoptions file idpd`. You can then set flags by entering `set security idp traceoptions flag all`.
 - The correct path for the idpd log file is `/var/log`, not `/var/db`.
- The “Configuring SSL Inspection” section of the *JUNOS Software Security Configuration Guide* incorrectly states that SSL inspection is disabled by default

and is enabled if any configurations are detected. This information is obsolete as of JUNOS Release 10.2. Although previously the SSL decoder was disabled by default in the detector, as of JUNOS Release 10.2 the SSL decoder is enabled by default. The updated information is as follows: The SSL decoder is enabled by default. To manually enable it, use the following CLI command:

```
user@host>set security idp sensor-configuration detector protocol-name
SSL tunable-name sc_ssl_flags tuneable-value 1
```

- The IDP rule notification options listed in the *JUNOS Software Security Configuration Guide* incorrectly include the **Send Emails** and **Run Scripts** options, which are not supported in the JUNOS 9.6 Release.

J-Web

- There is no documentation describing the J-Web pages for stateless firewall filters. To find these pages in J-Web, go to **Configuration > Firewall Filters**, then select **IPv4 Firewall Filters** or **IPv6 Firewall Filters**. After configuring filters, select **Assign to Interfaces** to assign your configured filters to interfaces.
- There is no documentation describing the J-Web pages for media gateways. To find these pages in J-Web, go to **Monitor > Media Gateway**.

Screens

- In the *Junos OS Design and Implementation Guide*, the “Implementing Firewall Deployments for Branch Offices” chapter contains incorrect screen configuration instructions.

Examples throughout this guide 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

- New Features in JUNOS Release 9.6 for J Series Services Routers on page 152
- Known Limitations in JUNOS Release 9.6 for J Series Services Routers on page 157
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series Services Routers on page 156
- Issues in JUNOS Release 9.6 for J Series Services Routers on page 159
- Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers on page 168
- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170

Hardware Requirements for JUNOS Release 9.6 for J Series Services Routers

- Transceiver Compatibility on page 168
- Power and Heat Dissipation Requirements for J Series PIMs on page 168
- Supported Third-Party Hardware for J Series Services Routers on page 168
- J Series CompactFlash and Memory Requirements on page 169

Transceiver Compatibility

We strongly recommend that only Juniper Networks transceivers be used on J Series interface modules. Different transceiver types (long-range, short-range, copper, and others) can be used together on multiport SFP interface modules as long as they are provided by Juniper Networks. We cannot guarantee that the interface module will operate correctly if third-party transceivers are used.

Please contact Juniper Networks for the correct transceiver part number for your device.

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 for J Series Services Routers

The following third-party hardware is supported for use with J Series Services Routers running JUNOS Software.

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 5 on page 169 lists the USB and CompactFlash card devices supported for use with the J Series Services Routers.

Table 5: 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 CompactFlash and Memory Requirements

Table 6 on page 169 lists the CompactFlash card and DRAM requirements for J Series Services Routers.

Table 6: 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**
- New Features in JUNOS Release 9.6 for J Series Services Routers on page 152
 - Known Limitations in JUNOS Release 9.6 for J Series Services Routers on page 157
 - Changes in Default Behavior and Syntax in JUNOS Release 9.6 for J Series Services Routers on page 156
 - Issues in JUNOS Release 9.6 for J Series Services Routers on page 159

- Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers on page 170
- Errata and Changes in Documentation for JUNOS Release 9.6 for J Series Services Routers on page 165

Upgrade and Downgrade Instructions for JUNOS Release 9.6 for J Series Services Routers

For upgrade and download instructions for JUNOS Software Release 9.6, please see the *JUNOS Software Migration Guide*.

JUNOS Software Release Notes for EX Series Switches

- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
- Limitations in JUNOS Release 9.6 for EX Series Switches on page 173
- Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
- Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
- Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179
- Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

New Features in JUNOS Release 9.6 for EX Series Switches

New features in Release 9.6 of JUNOS Software for EX Series switches are described in this section.

Not all EX Series software features are supported on all EX Series platforms in the current release. For a list of all EX Series software features and their platform support, see [EX Series Switch Software Features Overview](#).

New features are described on the following pages:

- Hardware on page 171
- Access Control and Port Security on page 171
- Ethernet Switching on page 171
- Layer 3 Protocols on page 171
- Management and RMON on page 172
- Packet Filters on page 172
- Virtual Chassis on page 172

Hardware

- **New optical transceiver support**—The SFP + uplink module in EX8200 switches now supports one new optical transceiver: EX-SFP-10GE-LRM (10GBase-LRM, 220 m).
- **New optical transceiver support**—The SFP uplink module in EX8200 switches now supports one new optical transceiver: EX-SFP-1GE-LX40K (1000Base-LX, 40 km).

Access Control and Port Security

- **Unrestricted proxy ARP**—For additional access port security on EX Series switches, you can use unrestricted proxy Address Resolution Protocol (ARP). With unrestricted proxy ARP, hosts cannot communicate directly with one another. Instead all communications must go through the switch. If you enable proxy ARP on an EX Series switch, the mode is unrestricted by default (that is the only mode supported) and proxy ARP applies globally to all interfaces on the switch whether you enabled it on one interface or more than one interface. The switch responds to any ARP request if the switch has an active route to the destination address.
- **Autorecovery from the disabled state on secure and storm control interfaces**—You can configure the switch to shut down interfaces due to a MAC limiting, MAC move limiting, or storm control error and automatically restore the disabled interfaces to service after a specified period of time. The shutdown option is a new option on storm control interfaces. The shutdown option for MAC limiting and MAC move limiting has been modified to include this temporary disabled state. You must configure the **port-error-disable** statement and specify a disable timeout to enable automatic restoration of the disabled interfaces to service.

Ethernet Switching

- **Extended Q-in-Q VLAN support**—Q-in-Q VLANs now support multiple S-VLANs per access interface, firewall-filter-based VLAN assignment, and routed VLAN interfaces (RVIs). Customer VLAN (C-VLAN) ranges now support rules for untagged and priority-tagged packets.

Layer 3 Protocols

- **IGMP snooping support for IGMPv3 INCLUDE mode**—IGMPv3 allows IGMP snooping to filter multicast streams based on the source address of the multicast stream. JUNOS Release 9.6 for EX Series switches supports IGMPv3 packets that are in INCLUDE mode. In INCLUDE mode, the switch requests that packets be sent to the specified multicast address only from specified IP source addresses.
- **Multicast VLAN registration**—Multicast VLAN registration (MVR) allows you to efficiently distribute IPTV multicast streams across an Ethernet ring-based Layer 2 network by creating a multicast VLAN (MVLAN), which becomes the only VLAN over which multicast traffic flows throughout the Layer 2 network. Multicast

traffic can then be selectively forwarded from ports on the MVLAN (source ports) to hosts that are connected to ports that are not part of the MVLAN.

Management and RMON

- **Automatic software download**—The automatic software download feature allows you to upgrade JUNOS Software on an EX Series switch as part of the DHCP process. The automatic software download feature works by your configuring the path to a software package on the DHCP server and the server communicating this information to DHCP clients, which are the switches. The DHCP clients that have been enabled for automatic software download receive these messages and, in cases where the software package in the DHCP server message is different from the software package that booted the DHCP client switch, download and install the software package.

Packet Filters

- **Extended support for firewall filter match conditions on EX8200 switches**—The following firewall filter match conditions are now added to the list of supported match conditions on EX8200 switches: `tcp-established`, `ip-options`, `source-prefix-list`, and `destination-prefix-list`.
- **Support for port-range as a match condition on EX3200 and EX4200 switches**—Packet Forwarding Engines on EX3200 and EX4200 switches now support the `port-range` match condition for firewall filters. To accommodate this new capability and provide the required hardware resources, support for the `packet-length` match condition in firewall filters has been removed starting at JUNOS Release 9.6.

Virtual Chassis

- **Link aggregation of uplink module ports configured as Virtual Chassis ports (VCPs)**—Link aggregation of SFP, SFP + , or XFP uplink module ports configured as VCPs across Virtual Chassis members is now supported in a Virtual Chassis configuration. When two or more pairs of such uplink VCPs are connected between two or more Virtual Chassis members and are running at the same link speeds, those VCPs are automatically formed into a link aggregation group (LAG).

Related Topics

- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
- Limitations in JUNOS Release 9.6 for EX Series Switches on page 173
- Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
- Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
- Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179
- Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches

- On EX3200 switches and EX4200 switches, the `request system power-off other-routing-engine` command and the `request system power-off both-routing-engines` command are disabled.
- Support for the `packet-length` match condition in firewall filters has been removed starting at JUNOS Release 9.6 for EX Series switches..

- Related Topics**
- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
 - Limitations in JUNOS Release 9.6 for EX Series Switches on page 173
 - Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
 - Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
 - Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179
 - Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

Limitations in JUNOS Release 9.6 for EX Series Switches

This section lists the limitations in JUNOS Release 9.6R4 for EX Series switches.

Class of Service

- On EX8200 switches, classification of packets using ingress firewall filter rules with forwarding-class and loss-priority configurations does not rewrite the DSCP or 802.1p bits. Rewriting of packets is determined by the forwarding-class and loss-priority values set in the DSCP classifier applied on the interface.

Infrastructure

- On EX Series switches, an SNMP query fails when the SNMP index size of a table is greater than 128 bytes, because the Net SNMP tool does not support SNMP index sizes greater than 128 bytes.
- Spanning tree, GVRP, or IGMP snooping configuration windows might load slowly in the J-Web interface. Wait till the windows load completely before entering information, or some information might get lost.

Interfaces

- EX Series switches do not support queued packet counters. Therefore, the queued packet counter in the output of the `show interfaces extensive` command always displays a count of 0 and is never updated.

- On EX3200 and EX4200 switches, when port mirroring is configured on all interfaces, the mirrored packets leaving a tagged interface might contain an incorrect VLAN ID.
- On EX8200 switches, port mirroring configuration on a Layer 3 interface with the output configured to a VLAN is not supported.

Related Topics

- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
- Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
- Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
- Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179
- Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

Outstanding Issues in JUNOS Release 9.6 for EX Series Switches

This section lists the outstanding issues in JUNOS Release 9.6R4 for EX Series switches.



NOTE: The following PRs that were previously included in the JUNOS Release 9.6 release notes as outstanding issues have been removed, because these issues are not present in JUNOS Release 9.6R4 for EX Series switches:

295588, 313195, 389276, 390812, 392043, 397290, 400360, 402163, 405899, 406032, 409321, 409934, 410947, 411660, 412908, 414110, 414213, 415748, 415959, 416062, 417024, 429589, 440611, 442373

The following are outstanding issues in JUNOS Release 9.6R4 for EX Series switches. The identifier following the description is the tracking number in our bug database.

Bridging, VLANs, and Spanning Trees

- On EX8200 switches, when the links on STP-enabled routed VLAN interfaces (RVIs) come up, control packets might egress before the STP BPDUs. [PR/300576]
- When Multiple VLAN Registration Protocol (MVRP) and MSTP are enabled together on EX Series switches, convergence does not occur between MVRP and MSTP. [PR/449248]
- On EX Series switches, when the VLAN with the lowest-numbered VLAN ID is down, the `show ntp associations` command output displays the following message:

```
/usr/bin/ntpq: write to localhost failed: No route to host
```

[PR/466595]

Class of Service

- On EX3200 and EX4200 switches, the **show interface queue** command output displays the count of transmitted packets and queued packets together under the field **Queued** instead of displaying the values under **Queued** and **Transmitted** fields. [PR/259525]
- On EX Series switches, if you deactivate one routing instance and the sub-interfaces associated with it, packets used by other routing instances that use sub-interfaces from the same main interface are classified incorrectly and sent to the wrong queues. [PR/493533]

Firewall Filters

- On EX Series switches, the **reject** action and the **log** or **syslog** action modifiers do not work as expected for packets destined for the Routing Engine. [PR/406714]

Hardware

- When an EX8216 switch power cycle completes, the **reason for last reboot** output on the master and backup Routing Engines might be displayed incorrectly. [PR/415569]
- On 48-port SFP line cards used in EX8200 switches, do not install a transceiver in the first or last port on the bottom row (ports 1 and 47). Transceivers installed in these ports are difficult to remove. As a workaround, you can remove the transceiver by using a small flathead screwdriver or other tool to lift the lock on the transceiver. [PR/423694]

Infrastructure

- 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]
- On EX8200 switches, RIP version 1 does not work properly. [PR/394905]
- In the J-Web interface, you cannot commit some configuration changes in the Port Configuration page and the VLAN Configuration page because of the following limitations for port-mirroring ports and port-mirroring VLANs:
 - A port configured as the output port for an analyzer cannot be a member of any VLAN other than the default VLAN.
 - A VLAN configured to receive analyzer output can be associated with only one port.

[PR/400814]
- When you issue the **request system power-off** command, the switch halts instead of turning off power. [PR/415772]

- In the J-Web interface, uploading a software package to the switch might not work properly if you are using Internet Explorer version 7. [PR/424859]
 - In the J-Web interface, the Ethernet Switching monitoring page might not display monitoring details if there are more than 13,000 MAC entries on the switch. [PR/425693]
 - If an SRE module, RE module, SF module, line card, or Virtual Chassis member is in offline mode, the J-Web interface might not update the dashboard image accordingly. [PR/431441].
 - In the J-Web interface, in the Port Security Configuration page, you are required to configure **action** when you configure **MAC limit** even though configuring an **action** value is not mandatory in the CLI. [PR/434836]
 - In the J-Web interface, interfaces configured with the **no-flow-control** statement might be displayed in the Link Aggregation Configuration page. [PR/437410]
 - In JUNOS Software for EX Series switches Releases 9.5 and 9.6, Layer 2 and Layer 3 traffic is VLAN-filtered in the Packet Forwarding Engine as the interface loses its VLAN membership when:
 - Graceful Routing Engine switchover (GRES) is performed.
 - Mastership switchover occurs.
 - The VCP cable is pulled out from the master switch or member switches.
 - The master switch or any member switch is rebooted.
 - The master switch or any member switch is halted.
 - Virtual Chassis split or merge is performed.
- [PR/438055]
- On EX3200 and EX4200 switches, in the Chassis Information page in the J-Web interface, the Fan tab under the Power and Fan Tray Details panel and the Temperature tab under the Chassis Component Details panel might display the FPC number prefixed to the names of the components. [PR/439264]
 - In the J-Web interface, changing port roles from Desktop, Desktop and Phone, and Layer 2 Uplink to other port roles might not remove the configurations for enabling dynamic ARP inspection and DHCP snooping. [PR/445080]
 - On EX Series switches, MAC addresses not present in the forwarding database (FDB) because of hash collision are not removed from the Ethernet switching process (**eswd**). These MAC addresses do not age out of the Ethernet switching table even if traffic is stopped completely and are never relearned when traffic is sent to these MAC addresses, even when there is no hash collision. As a workaround, clear those MAC addresses from the Ethernet switching table. [PR/451431]
 - On EX Series switches, aggregated Ethernet interfaces might go down when the software forwarding process (SFID) stops functioning and creates core files. [PR/452622]
 - In the J-Web interface, the DSCP classifiers associated with a logical interface might not appear to be mapped properly while you are editing the classifiers

associated with a logical interface. This issue might affect the Delete functionality also. [PR/455670]

- In the J-Web interface, the menu on the left side of the J-Web pages and contents of the J-Web pages might disappear when you double-click on the Troubleshoot tab. As a workaround, click on the Dashboard tab or Configure tab, and then click on the Troubleshoot tab to display the menu and contents of the page. [PR/459936]
- In a Virtual Chassis setup, the Chassis Information page (**Monitor > System View > Chassis Information**) in the J-Web interface displays an incorrect value for **Routing engine module** in the Master tab and no value for **Routing engine module** in the Backup tab. [PR/463811]
- J-Web interface creates session files even before the user logs in, and when a large number of J-Web sessions are invoked, Web management and other services become unavailable because of the large number of session files created in the `/var/jail/tmp/` directory. [PR/464897]
- If you attempt to set the time zone to Europe/Berlin on a switch with dual Routing Engines, the commit command might fail. [PR/483273]
- On EX Series switches, if you perform multiple commit checks and then commit the configuration, the CLI process might restart. [PR/485106]
- On EX4200 switches, under certain configurations (for example, when GRES is enabled and a backup member switch has a route to a destination whose egress member is on the backup member switch itself), packets generated on the backup member switch (for example, ping packets initiated from the backup switch to that destination) egress out of system. [PR/506119]

Interfaces

- The system log might display the following messages when the `monitor interfaces interface-name` command is issued simultaneously from multiple Telnet sessions:

```
Nov 21 11:55:29 8200-02-re0 /kernel: ifd_pfestat_req_wait_internal: ifd
ge-6/0/40, stats_req 0xa8f33d80, sreq_id 41028, new sreq_id 42053
Nov 21 11:55:44 8200-02-re0 login: LOGIN_INFORMATION: User regress logged
in from host 172.24.104.140 on device tty5
Nov 21 11:55:45 8200-02-re0 su: regress to root on /dev/tty5
Nov 21 11:55:53 8200-02-re0 /kernel: ifd_pfestat_req_wait_internal: ifd
ge-0/0/35, stats_req 0xa8a9dd20, sreq_id 4380, new sreq_id 5405
Nov 21 11:56:27 8200-02-re0 /kernel: ifd_pfestat_req_wait_internal: ifd
ge-0/0/30, stats_req 0xa8b60de0, sreq_id 54857, new sreq_id 55882
Nov 21 11:56:46 8200-02-re0 /kernel: ifd_pfestat_req_wait_internal: ifd
ge-0/0/31, stats_req 0xa89a56c0, sreq_id 36596, new sreq_id 37621
Nov 21 11:56:58 8200-02-re0 /kernel: ifd_pfestat_req_wait_internal: ifd
ge-0/0/33, stats_req 0xa8bd3d20, sreq_id 32622, new sreq_id 33647
Nov 21 11:57:08 8200-02-re0 /kernel: ifd_pfestat_req_wait_internal: ifd
ge-0/0/31, stats_req 0xa8bd3d20, sreq_id 52160, new sreq_id 53185
```

[PR/403842]

- On EX8200 switches, aggregated Ethernet interfaces might go down and come back up for a few minutes while the switch is updating many routes. [PR/416976]

- On EX8200 switches, when an egress VLAN that belongs to a routed VLAN interface (RVI) is configured as the input for a port mirroring analyzer, the analyzer incorrectly appends a dot1q (802.1Q) header to the mirrored packets or does not mirror any packets at all. As a workaround, configure a port mirroring analyzer with each port of the VLAN as egress input. [PR/445393]

Layer 2 and Layer 3 Protocols

- IGMP snooping does not function for IGMPv3 reports with the exclude filter mode. [PR/286600]

Related Topics

- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
- Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
- Limitations in JUNOS Release 9.6 for EX Series Switches on page 173
- Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
- Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179
- Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

Resolved Issues in JUNOS Release 9.6 for EX Series Switches

The following are the issues that have been resolved since JUNOS Release 9.6R4 for EX Series switches. The identifier following the descriptions is the tracking number in our bug database.

Access Control and Port Security

- On EX Series switches, if you configure the RADIUS server `revert-interval interval` option, the switch does not attempt to reconnect to the unreachable server after the revert interval has elapsed. [PR/304637: This issue has been resolved.]

Infrastructure

- In the J-Web interface, the Add button in the IGMP Snooping Configuration page might be disabled even when VLANs are configured on the switch. [PR/460157: This issue has been resolved.]
- On EX Series switches, the /var directory appears full after some files in the /var/log directory are deleted. To avoid this problem, use the clear log filename command to clear the log files, instead of deleting them manually. [PR/496298: This issue has been resolved.]

- Related Topics**
- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
 - Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
 - Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
 - Limitations in JUNOS Release 9.6 for EX Series Switches on page 173
 - Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179
 - Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

Errata in Documentation for JUNOS Release 9.6 for EX Series Switches

This section lists outstanding issues with the documentation.

There are no errata at this release.

- Related Topics**
- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
 - Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
 - Limitations in JUNOS Release 9.6 for EX Series Switches on page 173
 - Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
 - Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
 - Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches on page 179

Upgrade and Downgrade Issues for JUNOS Release 9.6 for EX Series Switches

The following pages list the issues in JUNOS Release 9.6R4 for EX Series switches regarding software upgrade or downgrade:

Upgrading or Downgrading from JUNOS Release 9.4R1 for EX Series Switches

The ARP aging time configuration in the **system** configuration stanza in JUNOS Release 9.4R1 is incompatible with the ARP aging time configuration in JUNOS Release 9.3R1 or earlier and JUNOS Release 9.4R2 or later. If you have configured **system arp aging-timer aging-time** on EX Series switches running JUNOS Release 9.4R1 and upgrade to JUNOS Release 9.4R2 or later or downgrade to JUNOS Release 9.3R1 or earlier, the switch will display configuration errors on booting up after the upgrade or downgrade. As a workaround, delete the **arp aging-timer aging-time** configuration in the **system** configuration stanza and reapply the configuration after you complete the upgrade or downgrade.

The format of the file in which the Virtual Chassis topology information is stored has changed in JUNOS Release 9.4. When you downgrade JUNOS Release 9.4 or later running on EX4200 switches in a Virtual Chassis to JUNOS Release 9.3 or earlier, make topology changes, and then upgrade to JUNOS Release 9.4 or later, the topology changes you have made using JUNOS Release 9.3 or earlier are not retained. The switch restores the last topology change you have made using JUNOS Release 9.4.

Upgrading from JUNOS Release 9.3 to Release 9.6 for EX Series Switches

If you are upgrading from JUNOS Release 9.3R1 and have voice over IP (VoIP) enabled on a private VLAN (PVLAN), you must remove this configuration before upgrading, to prevent upgrade problems. VoIP on PVLAN ports is not supported on releases after JUNOS Release 9.3R1.

Upgrading from JUNOS Release 9.2 to Release 9.6 for EX Series Switches

For JUNOS Release 9.3 and later 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-suplicant 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 9.6. 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 is not configured and you try to commit the configuration, the commit will fail.
- Remove static MAC addresses on single-suplicant mode interfaces. If they exist and you try to commit the configuration, the commit will fail.
- 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 supported in a static MAC configuration. If they exist and you try to commit the configuration, the commit will fail.
- Support for static MAC bypass in single or single-secure mode has been removed. If static MAC bypass exists and you try to commit the configuration, the commit will fail.
- In an 802.1X configuration stanza, the switch will not accept the option `vrange` as an assigned VLAN name. If it exists and you try to commit the configuration, the commit will fail.
- Enabling 802.1X and the port mirroring feature on the same interface is not supported. If you enable 802.1X and port mirroring 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.
- If the MSTP configuration contains a VLAN (under `protocols mstp msti msti-id`) that does not exist on the switch and you try to commit the configuration, the commit will fail. Remove the VLAN from the MSTP configuration before you perform an upgrade.
- In the `interfaces` configuration stanza, if `no-auto-negotiation` is configured but speed and link duplex settings are not configured under `ether-options` and you try to commit the configuration, the commit will fail. If `no-auto-negotiation` is configured under `ether-options`, you must configure speed and link duplex settings.
- In the `ethernet-switching-options` configuration, if `action` is not configured for the number of MAC addresses allowed on the interface (under `secure-access-port interface interface-name mac-limit` in the CLI or in the Port Security Configuration page in the J-Web interface), and you try to commit the configuration, the commit will fail. You must configure an action for the MAC address limit before upgrading from Release 9.2 to Release 9.6.
- If you have configured a tagged interface on logical interface 0 (unit 0), configure a tagged interface on a logical interface other than unit 0 before upgrading from Release 9.2 to Release 9.6. If you have not done this and you try to commit the configuration, the commit will fail. Beginning with JUNOS Release 9.3 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.
- On EX4200 switches, if you have installed advanced licenses for features such as BGP, rename the `/config/license` directory to `/config/.license_priv` before upgrading from Release 9.2 to Release 9.3 or later. If the switch does not have a `/config/license` directory, create the `/config/.license_priv` directory manually before you upgrade. If you do not rename the `/config/license` directory or create the `/config/.license_priv` directory manually, the licenses installed will be deleted after you upgrade from Release 9.2 to Release 9.3 or later.

Downgrading from JUNOS Release 9.6 to Release 9.2 for EX4200 Switches

When you downgrade a Virtual Chassis configuration from JUNOS Release 9.6 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.

- Related Topics**
- New Features in JUNOS Release 9.6 for EX Series Switches on page 170
 - Changes in Default Behavior and Syntax in JUNOS Release 9.6 for EX Series Switches on page 173
 - Limitations in JUNOS Release 9.6 for EX Series Switches on page 173

- Outstanding Issues in JUNOS Release 9.6 for EX Series Switches on page 174
- Resolved Issues in JUNOS Release 9.6 for EX Series Switches on page 178
- Errata in Documentation for JUNOS Release 9.6 for EX Series Switches on page 179

JUNOS Documentation and Release Notes

For a list of related JunosE documentation, see <http://www.juniper.net/techpubs/software/junos/>.

If the information in the latest release notes differs from the information in the documentation, follow the *JunosE Release Notes*.

To obtain the most current version of all Juniper Networks® technical documentation, see the product documentation page on the Juniper Networks website at <http://www.juniper.net/techpubs/>.

Juniper Networks supports a technical book program to publish books by Juniper Networks engineers and subject matter experts with book publishers around the world. These books go beyond the technical documentation to explore the nuances of network architecture, deployment, and administration using the Junos operating system (Junos OS) and Juniper Networks devices. In addition, the Juniper Networks Technical Library, published in conjunction with O'Reilly Media, explores improving network security, reliability, and availability using Junos OS configuration techniques. All the books are for sale at technical bookstores and book outlets around the world. The current list can be viewed at <http://www.juniper.net/books>.

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

01 June 2010—Revision 4, JUNOS Release 9.6R4
01 February 2010—Revision 3, JUNOS Release 9.6R3
08 October 2009—Revision 2, JUNOS Release 9.6R2
05 August 2009—Revision 1, JUNOS Release 9.6R1

Copyright © 2010, Juniper Networks, Inc. All rights reserved.

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

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

Products made or sold by Juniper Networks or components thereof might be covered by one or more of the following patents that are owned by or licensed to Juniper Networks: U.S. Patent Nos. 5,473,599, 5,905,725, 5,909,440, 6,192,051, 6,333,650, 6,359,479, 6,406,312, 6,429,706, 6,459,579, 6,493,347, 6,538,518, 6,538,899, 6,552,918, 6,567,902, 6,578,186, and 6,590,785.