

# Junos<sup>®</sup> OS 12.2 Release Notes

Release 12.2R9  
5 August  
Revision 4

These release notes accompany Release 12.2R9 of the Junos operating system (Junos OS). They describe device documentation and known problems with the software. Junos OS runs on all Juniper Networks M Series, MX Series, and T Series routing platforms, EX Series Ethernet Switches, and the ACX Series.

For the latest, most complete information about outstanding and resolved issues with the Junos OS software, see the Juniper Networks online software defect search application at <http://prsearch.juniper.net>.

You can also find these release notes on the Juniper Networks Junos OS Documentation Web page, which is located at <https://www.juniper.net/techpubs/software/junos/>.

## Contents

Junos OS Release Notes for ACX Series Routers . . . . .	5
New Features in Junos OS Release 12.2 for ACX Series Routers . . . . .	5
Hardware . . . . .	6
Class of Service (CoS) . . . . .	6
Infrastructure . . . . .	7
Interfaces and Chassis . . . . .	7
Layer 2 and Layer 3 Protocols . . . . .	14
MPLS . . . . .	17
Network Management and Monitoring . . . . .	19
Power Management . . . . .	19
Routing Policy and Firewall Filters . . . . .	20
Routing Protocols . . . . .	21
Software Architecture . . . . .	22
Timing and Synchronization . . . . .	23
Changes in Default Behavior and Syntax in Junos OS Release 12.2 for ACX Series Routers . . . . .	29
IPv6 . . . . .	29
Known Limitations in Junos OS Release 12.2 for ACX Series Routers . . . . .	29
Class of Service (CoS) . . . . .	29
Interfaces and Chassis . . . . .	29
Multiprotocol Label Switching (MPLS) . . . . .	30
Routing Policy and Firewall Filters . . . . .	30

Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers . . . . .	31
Interfaces and Chassis . . . . .	31
Resolved Issues in Junos OS Release 12.2 for ACX Series Routers . . . . .	31
Current Release . . . . .	31
Previous Releases . . . . .	31
Errata and Changes in Documentation for Junos OS Release 12.2 for ACX	
Series Routers . . . . .	34
Hardware . . . . .	34
Class of Service (CoS) . . . . .	34
Routing Protocols . . . . .	36
Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX	
Series Routers . . . . .	36
Basic Procedure for Upgrading to Release 12.2 . . . . .	36
Upgrade and Downgrade Support Policy for Junos OS Releases . . . . .	39
Junos OS Release Notes for EX Series Switches . . . . .	41
New Features in Junos OS Release 12.2 for EX Series Switches . . . . .	41
Hardware . . . . .	42
Access Control and Port Security . . . . .	43
Ethernet Switching and Spanning Trees . . . . .	44
Firewall Filters . . . . .	45
High Availability . . . . .	45
Infrastructure . . . . .	46
Interfaces . . . . .	46
J-Web Interface . . . . .	47
Management and RMON . . . . .	47
MPLS . . . . .	47
Power over Ethernet (PoE) . . . . .	48
Software Installation and Upgrade . . . . .	48
Virtual Chassis . . . . .	48
Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX	
Series Switches . . . . .	49
Ethernet Switching and Spanning Trees . . . . .	49
High Availability . . . . .	49
Infrastructure . . . . .	49
Limitations in Junos OS Release 12.2 for EX Series Switches . . . . .	50
Access Control and Port Security . . . . .	50
Ethernet Switching and Spanning Trees . . . . .	51
Firewall Filters . . . . .	51
Hardware . . . . .	51
High Availability . . . . .	52
Infrastructure . . . . .	52
Interfaces . . . . .	53
J-Web Interface . . . . .	54
Layer 2 and Layer 3 Protocols . . . . .	55
Management and RMON . . . . .	55
Software Installation and Upgrade . . . . .	56
Virtual Chassis . . . . .	56

Outstanding Issues in Junos OS Release 12.2 for EX Series Switches . . . . .	57
Access Control and Port Security . . . . .	58
Hardware . . . . .	58
Infrastructure . . . . .	58
Interfaces . . . . .	60
J-Web Interface . . . . .	60
Routing Protocols . . . . .	63
Software Upgrade and Installation . . . . .	63
Spanning-Tree Protocols . . . . .	63
Virtual Chassis . . . . .	63
Resolved Issues in Junos OS Release 12.2 for EX Series Switches . . . . .	64
Issues Resolved in Release 12.2R1 . . . . .	64
Issues Resolved in Release 12.2R2 . . . . .	73
Issues Resolved in Release 12.2R3 . . . . .	77
Issues Resolved in Release 12.2R4 . . . . .	78
Issues Resolved in Release 12.2R5 . . . . .	80
Issues Resolved in Release 12.2R6 . . . . .	83
Issues Resolved in Release 12.2R7 . . . . .	85
Issues Resolved in Release 12.2R8 . . . . .	87
Issues Resolved in Release 12.2R9 . . . . .	89
Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches . . . . .	92
Changes to Junos OS for EX Series Switches Documentation . . . . .	92
Errata . . . . .	92
Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches . . . . .	94
Upgrade and Downgrade Support Policy for Junos OS Releases . . . . .	94
Upgrading or Downgrading to Junos OS Release 12.2R1 . . . . .	94
Upgrading to Junos OS Release 12.1R2 or Later, with Existing VSTP Configurations . . . . .	95
Upgrading from Junos OS Release 10.4R3 or Later . . . . .	95
Upgrading from Junos OS Release 10.4R2 or Earlier . . . . .	96
Upgrading EX Series Switches Using NSSU . . . . .	97
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers . . . . .	100
New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers . . . . .	100
Class of Service (CoS) . . . . .	101
Forwarding and Sampling . . . . .	107
High Availability (HA) and Resiliency . . . . .	108
Interfaces and Chassis . . . . .	110
Junos OS Installation and Upgrade . . . . .	127
Junos OS XML API and Scripting . . . . .	128
Layer 2 Ethernet Services . . . . .	129
Multiprotocol Label Switching (MPLS) . . . . .	133
Multicast . . . . .	134
Network Management and Monitoring . . . . .	137
Routing Policy and Firewall Filters . . . . .	138
Routing Protocols . . . . .	140

Services Applications . . . . .	140
Subscriber Access Management . . . . .	140
User Interface and Configuration . . . . .	150
VPNs . . . . .	150
Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M	
Series, MX Series, and T Series Routers . . . . .	154
Changes in Default Behavior and Syntax . . . . .	154
Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T	
Series Routers . . . . .	170
Class of Service (CoS) . . . . .	170
Subscriber Management and Services . . . . .	171
Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series	
Routers . . . . .	171
Current Software Release . . . . .	172
Previous Releases . . . . .	201
Errata and Changes in Documentation for Junos OS Release 12.2 for M Series,	
MX Series, and T Series Routers . . . . .	285
Errata . . . . .	285
Changes to the Junos OS Documentation Set . . . . .	321
Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series,	
MX Series, and T Series Routers . . . . .	322
Basic Procedure for Upgrading to Release 12.2 . . . . .	322
Upgrade and Downgrade Support Policy for Junos OS Releases . . . . .	325
Upgrading a Router with Redundant Routing Engines . . . . .	326
Upgrading Juniper Network Routers Running Draft-Rosen Multicast	
VPN to Junos OS Release 10.1 . . . . .	326
Upgrading the Software for a Routing Matrix . . . . .	328
Upgrading Using Unified ISSU . . . . .	329
Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled	
for Both PIM and NSR . . . . .	329
Downgrading from Release 12.2 . . . . .	330
Junos OS Documentation and Release Notes . . . . .	332
Documentation Feedback . . . . .	332
Requesting Technical Support . . . . .	332
Revision History . . . . .	334

## Junos OS Release Notes for ACX Series Routers

---

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Resolved Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers on page 34](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

### New Features in Junos OS Release 12.2 for ACX Series Routers

Powered by Junos OS, ACX Series Universal Access Routers provide superior management for rapid provisioning to the access network. They are designed to support residential, mobile, and business access. ACX Series routers include the ACX1000, ACX1100, ACX2000, and the ACX2100 routers.

The following are key features of ACX Series routers:

- High performance up to 10 Gigabit Ethernet capable
- Seamless MPLS traffic engineering for optimal paths and per-customer quality of service in the access layer
- Built-in Precision Timing Protocol (PTP) and Synchronized Ethernet to eliminate dropped calls and data retransmissions
- Environmentally hardened with 65 W Power over Ethernet Plus (PoE+)

The following features have been added to Junos OS Release 12.2 for ACX Series Universal Access Routers. Following the description is the title of the manual or manuals to consult for further information:

- [Hardware on page 6](#)
- [Class of Service \(CoS\) on page 6](#)
- [Infrastructure on page 7](#)
- [Interfaces and Chassis on page 7](#)
- [Layer 2 and Layer 3 Protocols on page 14](#)
- [MPLS on page 17](#)
- [Network Management and Monitoring on page 19](#)
- [Power Management on page 19](#)
- [Routing Policy and Firewall Filters on page 20](#)
- [Routing Protocols on page 21](#)

- [Software Architecture on page 22](#)
- [Timing and Synchronization on page 23](#)

## Hardware

---

- **New ACX1000 Universal Access Router**—Starting in Release 12.2, Junos OS supports the ACX1000 router. This router enables a wide range of business and residential applications and services, including microwave cell site aggregation, MSO mobile backhaul service cell site deployment, and service provider or operator cell site deployment. The ACX1000 router is a compact access router that is one rack unit (U) tall. The ACX1000 router contains 8 T1 and E1 ports and 8 Gigabit Ethernet ports. The ACX1000 router also supports either 4 RJ45 (Cu) ports or installation of 4 Gigabit Ethernet SFP transceivers.

[See [ACX1000 Universal Access Router](#).]

- **New ACX2000 Universal Access Router**—Starting in Release 12.2, Junos OS supports the ACX2000 router. This router enables a wide range of business and residential applications and services, including microwave cell site aggregation, MSO mobile backhaul service cell site deployment, and service provider or operator cell site deployment. The ACX2000 router is a compact access router that is one rack unit (U) tall. The ACX2000 router contains 16 T1 and E1 ports, 6 Gigabit Ethernet ports, and 2 PoE ports. The ACX2000 router also supports installation of two Gigabit Ethernet SFP transceivers and two 10-Gigabit Ethernet SFP+ transceivers.

[See [ACX2000 Universal Access Router](#).]

## Class of Service (CoS)

---

- **Existing CoS features supported on the ACX Series Universal Access Routers**—Existing Junos OS class-of-service (CoS) features are supported without changes to statements or functionality.

The following key CoS features are supported:

- Physical interface-based classifiers at the **[edit class-of-service interfaces *interfaces-name*]** hierarchy level.
- Fixed classification for all ingress packets traversing a logical interface to a single forwarding class. Fixed classification is supported on all interfaces types.
- Experimental (EXP) bits located in each MPLS label and used to encode the CoS value of a packet as it traverses an LSP. To configure global EXP bits, include the **exp** statement at the **[edit class-of-service system-defaults classifiers]** hierarchy level.
- Attachment of the following rewrite rules to the physical interface at the **[edit class-of-service interfaces *interface-name* rewrite-rules]** hierarchy level: IP ToS, DSCP, and IEEE 802.1p bit value.
- Rewrite rules for MPLS EXP bits on the logical interface at the **[edit class-of-service interfaces *interface-name* unit *unit-number* rewrite-rule]** hierarchy level.



**NOTE:** Fine-grained rewrite is not possible, even when you use multifield filters.

Queuing and scheduling features include:

- Support for up to eight forwarding classes.
- Up to eight egress queues per port.
- Internal buffer of 2 MB with per-egress queue buffer management.
- Three weighted random early detection (WRED) curves for TCP and one WRED curve for non-TCP. There are two fill levels and two drop probabilities per WRED curve; the drop probability corresponding to the first fill must be zero.
- Strict-priority and weighted deficit round-robin scheduling.
- Multiple strict-priority queues per port.
- Per-queue committed information rate (CIR) and peak information rate (PIR).
- Per-physical-port shaping.

Queue statistics features include:

- Per-egress-queue enqueue statistics in packets, bytes, packets per second (pps), and bits per second (bps).
- Per-egress-queue transmit statistics in packets, bytes, pps, and bps.
- Per-egress-queue drop statistics in packets and pps.

## Infrastructure

- **Dual-root partitioning**—All ACX Series routers support dual-root partitioning. Dual-root partitioning means that the primary and backup Junos OS images are kept in two independently bootable root partitions. If the primary root partition becomes corrupted, the system remains fully functional by booting from the backup Junos OS image located in the other root partition.

[See [Dual-Root Partitioning ACX Series Universal Access Routers Overview](#).]

## Interfaces and Chassis

- **Junos OS support for chassis management of ACX Series Universal Access Routers**—Junos OS Release 12.2 supports the following ACX Series routers:
  - ACX1000 Universal Access Router
  - ACX2000 Universal Access Router

The ACX Series router chassis does not have redundancy support.

The following CLI operational mode commands support chassis management operations on an ACX Series router:

Show commands:

- **show chassis alarms**
- **show chassis craft-interface**
- **show chassis environment**
- **show chassis feb**
- **show chassis firmware**
- **show chassis fpc < pic-status >**
- **show chassis hardware < clei-models | detail | extensive | models >**
- **show chassis mac-addresses**
- **show chassis routing-engine**
- **show chassis pic fpc-slot *fpc-slot* pic-slot *pic slot***

Request command:

- **request chassis feb restart slot *slot-number***

Restart command:

- **restart chassis-control < gracefully | immediately | soft >**

[See [System Basics: Chassis-Level Features Configuration Guide](#).]

- **Gigabit Ethernet physical interface features (ACX Series Universal Access Routers)**—The following Gigabit Ethernet physical interface features are supported on ACX Series Universal Access routers:
  - **Autonegotiation for Gigabit Ethernet interfaces**—Exchange of the following parameters is supported: speed and duplex mode. Autonegotiation can be enabled or disabled. When autonegotiation is disabled, the speed has to be explicitly configured to 10–100 Mbps. To configure autonegotiation, include the **auto-negotiation** statement at the [edit interfaces *interface-name* **gigether-options**] hierarchy level. To disable the autonegotiation, include the **no-auto-negotiation** statement at the [edit interfaces *interface-name* **gigether-options**] hierarchy level.  
[See [Gigabit Ethernet Autonegotiation Overview](#) and [Junos OS Ethernet Interfaces Configuration Guide](#).]
  - **Event handling of SFP insertion and removal**—When you insert a small form-factor pluggable transceiver (SFP), the port needs to be configured with the correct speed for that interface (Gigabit Ethernet or 10-Gigabit Ethernet). The following details apply to SFP insertion and removal:
    - SFP-based 1-Gigabit Ethernet interfaces support the following standards:
      - 1000BASE-SX
      - 1000BASE-LX



- 1000BASE-T
- 100BASE-FX (100M)
- The 10-Gigabit Ethernet interfaces based on SFP+ support the following standards in addition to the 1-Gigabit Ethernet interface standards mentioned above:
  - 10GBASE-SR
  - 10GBASE-LR
- On an SFP+ port, the port speed is not set by autonegotiation. Instead, it is determined by the speed of the SFP that is inserted or removed. The default speed of the SFP+ port is 10 Gbps. However, when a Gigabit Ethernet SFP is inserted in the SFP+ slot, Junos OS changes the speed to 1 Gbps. When the Gigabit Ethernet SFP is removed, the port speed is automatically reset to the default 10 Gbps.

[See [Junos OS Interfaces Fundamentals Configuration Guide](#).]

- **Explicit disabling of the physical interface**—Disable a physical interface by effectively unconfiguring it. To disable an interface, include the **disable** statement at the [**edit interfaces interface-name**] hierarchy level.

[See [disable \(Interface\)](#).]

- **Loopback**—Local loopback is supported at the **gigether-options** hierarchy level. Local loopback allows packets to flow in toward the system. To configure the local loopback, include the **loopback** statement at the [**edit interfaces interface-name gigether-options**] hierarchy level.

[See [loopback \(Aggregated Ethernet, Fast Ethernet, and Gigabit Ethernet\)](#).]

- **Loss of signal (LOS) alarm**—A LOS alarm indicates that a signal could not be detected at the physical interface level. The LOS alarm is generated by the physical interface and displays a Link Up or Link Down event. To display LOS and other alarms, issue the **show interfaces interface-name extensive** command.

[See [show interfaces extensive](#).]

- **Maximum transmission unit (MTU)**—Specify the MTU size for the interface. To configure the MTU, specify the **bytes** in the **mtu** statement at the [**edit interfaces interface-name**] hierarchy level.

[See [Configuring the Media MTU](#).]

- **Remote fault notification for 10-Gigabit Ethernet interfaces**—Notifies each end of a connection of the failure at that end. When the failure is identified, the link is brought down and the LED light is turned off. This feature is not user configured.

[See [Detecting Remote Faults](#).]

- **Statistics collection and handling**—Port-level input and output error statistics and the logical interface level statistics are collected automatically from the Packet Forwarding Engine. To display statistics, issue the **show interfaces interface-name (brief | extensive)** operational mode command.

[See [show interfaces statistics](#).]



**NOTE:** The ACX Series routers do not support flow control based on PAUSE frames.

[See [Junos OS Ethernet Interfaces Configuration Guide](#) and [Junos OS System Basics Configuration Guide](#).]

- **Media type selection (ACX1000 Universal Access routers)**—You can select the media type (copper or fiber) for the 1-Gigabit Ethernet interfaces. To specify the media type, include the new **media-type** statement with the **copper** or **fiber** option at the [edit **interfaces interface-name**] hierarchy level.



**NOTE:** Media type selection is applicable to ports only in slot 2.

[See [Junos OS Ethernet Interfaces Configuration Guide](#).]

- **IEEE 802.1ag OAM CFM and ITU-T Y.1731**—The ACX Series routers support the IEEE 802.1ag standard for Operation, Administration, and Management (OAM) connectivity fault management (CFM) and the ITU-T Y.1731 standard for Ethernet service OAM.

The IEEE 802.1ag standard defines mechanisms for end-to-end Ethernet service assurance over any path, whether a single link or multiple links spanning networks composed of multiple LANs.

The ITU-T Y.1731 uses different terminology than IEEE 802.1ag and in addition defines Ethernet service OAM features for fault monitoring, diagnostics, and performance monitoring.

The following key CFM and Ethernet service OAM features are supported:

- Continuity check
- Loopback messages
- Traceroute messages
- Linktrace messages

In addition, the following key ITU-T Y.1731 Ethernet Service OAM features are supported:

- Performance monitoring
  - Delay measurements
  - Loss measurements



**NOTE:** Maintenance association intermediate points (MIPs) are not supported on the ACX Series routers.



**NOTE:** The test signal, automatic protection switching, maintenance communication channel, experimental, and vendor-specific PDUs are not supported for generation or receipt in Junos OS or on the ACX Series routers.

The proactive and dual-ended loss measurement functionality of ITU-T Y1731 is not supported.

[See [IEEE 802.1ag OAM Connectivity Fault Management Overview](#), [ITU-T Y.1731 Ethernet Service OAM](#), and [Ethernet Interfaces](#).]

- **IEEE 802.3ah OAM link-fault management**—The ACX Series routers support the IEEE 802.3ah standard for Operation, Administration, and Maintenance (OAM). The IEEE 802.3ah standard defines a set of link fault management mechanisms to detect and report link faults on a single point-to-point Ethernet LAN. The following OAM link fault management features are supported:

- Discovery
- Link monitoring
- Remote fault detection
- Remote loopback

[See [IEEE 802.3ah OAM Link-Fault Management Overview](#).]

- **Junos OS support for ACX Series Universal Access Routers**—Starting with Release 12.2R2, Junos OS supports the following ACX Series routers:

- ACX1100 Universal Access Routers
- ACX2100 Universal Access Routers

The following **show chassis** commands are supported on the ACX1100 and the ACX2100 routers:

- **show chassis fpc**
- **show chassis fpc fpc-slot**
  - On ACX1100 routers, replace **fpc-slot** with value **0**.
  - On ACX2100 routers, replace **fpc-slot** with a value **0** through **1**.
- **show chassis fpc detail**
- **show chassis fpc pic-status**

[See [show chassis fpc](#).]

- **T1 and E1 interfaces time-division multiplexing (TDM) support**—Existing Junos OS TDM features are supported without changes to statements or functionality. The following key TDM features for Channelized T1 (**ct1**) interfaces and Channelized E1 (**ce1**) interfaces are supported:

- **T1/E1 ports**—The ACX1000 router has 8 built-in TDM ports. The ACX2000 router has 16 built-in TDM ports. T1/E1 mode selection is at the PIC level. To set the T1/E1 mode, include the **framing** statement with the **t1** or **e1** option at the [**chassis fpc 0 pic slot-number**] hierarchy level. All ports can be T1 or E1. Mixing T1s and E1s is not supported.

[See [framing](#).]

- **T1/E1 channelization**—Full channelization is supported. Partitioning is not supported. To configure full channelization, include the **no-partition** statement at the [**edit interfaces ct1-fpc/pic/port**] hierarchy level or at the [**edit interfaces ce1-fpc/pic/port**] hierarchy level, depending on the interface type.

[See [no-partition](#).]

- **T1/E1 encapsulation**—Structure-Agnostic TDM over Packet (SAToP) defined in RFC 4553 is supported. SAToP is used to transport complete TDM frames across the transport network, creating a smooth migration from legacy TDM to the central office. Traffic is kept at a constant bit rate of 1.544 Mbps for T1 and 2.048 Mbps plus overhead for E1 interfaces.

[See [SAToP Emulation on T1 and E1 Interfaces Overview](#).]

- **Alarms, defects, and statistics**—Display alarms, defects, and statistics for interfaces running on the ACX Series routers.

[See [show interfaces \(T1 or E1\)](#).]

- **BERT algorithms**—Run BERT for interfaces running on the ACX Series routers.

[See [Configuring T1 BERT Properties](#) and [test interface t1-bert-start](#).]

- **External and internal loopback**—Use loopback testing to isolate interface problems. By default, loopback is not configured.

[See [Configuring T1 Loopback Capability](#), [Configuring E1 Loopback Capability](#), [Junos OS Interfaces Network Operations Guide](#), and [Junos OS E1/E3/T1/T3 Interfaces Configuration Guide](#).]

- **ATM time-division multiplexing (TDM) support**—Existing Junos OS TDM features are supported without changes to statements or functionality. The following key TDM features for ATM are supported:

- **Inverse multiplexing for ATM (IMA)**—Defined by the ATM Forum IMA specification version 1.1. IMA is a standardized technology used to transport ATM traffic over a bundle of T1 and E1 interfaces, also known as an IMA group. Up to 8 links per bundle and 16 bundles per PIC are supported.

[See [Configuring Inverse Multiplexing for ATM \(IMA\)](#).]

- **Inverse multiplexing for ATM (IMA) Layer 2 encapsulation**—Layer 2 encapsulation for IMA pseudowire initiation and termination on the ACX Series routers is supported. To configure encapsulation at the logical interface level, include the **encapsulation** statement with the **atm-ccc-cell-relay** or **atm-ccc-vc-mux** option at the [**edit interface interface-name unit logical-unit-number**] hierarchy level.

[See [Understanding Encapsulation on an Interface \(ACX Series Routers\)](#).]

- **Denied packets counter**—The **show interfaces** command for ATM interfaces, **show interfaces at-fpc/pic/port extensive**, supports a new field: **denied packets**. The **denied packets** field displays the number of packets dropped because of VLAN priority deny packets or because of an error in the forwarding configuration that might cause a negative frame length, that is, the stripping size is larger than the packet size.

[See [show interfaces \(ATM\)](#).]

- **TDM and ATM class-of-service (CoS)**—Junos OS CoS enables you to classify traffic into classes and offer various levels of throughput and packet loss when congestion occurs. Fixed classification is supported on the ACX Series routers. To configure fixed classification, include the **forwarding-class** statement at the [**edit class-of-service interfaces interface-name unit logical-unit-number**] hierarchy level.

[See [forwarding-class \(Interfaces\)](#) and [CoS on ACX Series Universal Access Routers Features Overview](#).]

- **ATM policing and shaping**—Policing, or rate limiting, is an important component of firewall filters that you can use to limit the amount of traffic that passes into or out of an interface. Shaping uses queuing and scheduling to shape the outgoing traffic. For more information about supported policing and shaping on the ACX Series routers, see the Firewalls section of these release notes.

[See [Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview](#).]

- **TDM CESoPSN (ACX1000 and ACX2000 routers)**—Structure-aware TDM Circuit Emulation Service over Packet Switched Network (CESoPSN) is a method of encapsulating TDM signals into CESoPSN packets, and in the reverse direction, decapsulating CESoPSN packets back into TDM signals—also, referred to as Interworking Function (IWF). The following CESoPSN features are supported:
  - **Channelization up to the ds0 level**—The following numbers of NxDS0 pseudowires are supported for 16 T1 and E1 built-in ports and 8 T1 and E1 built-in ports.  
Sixteen T1 and E1 built-in ports support the following number of pseudowires:
    - Each T1 port can have up to 24 NxDS0 pseudowires, which add up to a total of up to 384 NxDS0 pseudowires.
    - Each E1 port can have up to 31 NxDS0 pseudowires, which add up to a total of up to 496 NxDS0 pseudowires.
 Eight T1 and E1 built-in ports support the following number of pseudowires:
    - Each T1 port can have up to 24 NxDS0 pseudowires, which add up to a total of up to 192 NxDS0 pseudowires.
    - Each E1 port can have up to 31 NxDS0 pseudowires, which add up to a total of up to 248 NxDS0 pseudowires.
  - **Protocol support**—All protocols that support Structure-Agnostic TDM over Packet (SAToP) support CESoPSN NxDS0 interfaces.
  - **Packet latency**—The time required to create packets (from 1000 through 8000 microseconds).

- **CESoPSN encapsulation**—The following statements are supported at the [edit interfaces *interface-name*] hierarchy level:
  - *ct1-x/y/z* partition *partition-number* timeslots *timeslots* interface-type *ds*
  - *ds-x/y/z:n* encapsulation *cesopsn*
- **CESoPSN options**—The following statements are supported at the [edit interfaces *interface-name* *cesopsn-options*] hierarchy level:
  - *excessive-packet-loss-rate* (sample-period *milliseconds*)
  - *idle-pattern* *pattern*
  - *jitter-buffer-latency* *milliseconds*
  - *jitter-buffer-packets* *packets*
  - *packetization-latency* *microseconds*
- **Interfaces show commands**—The *show interfaces interface-name extensive* command is supported for *t1*, *e1*, and *at* interfaces.
- **CESoPSN pseudowires**—CESoPSN pseudowires are configured on the logical interface, not on the physical interface. So the *unit logical-unit-number* statement must be included in the configuration at the [edit interfaces *interface-name*] hierarchy level. When you include the *unit logical-unit-number* statement, circuit cross-connect (CCC) for the logical interface is created automatically.

[ *ACX Series Universal Access Router Configuration Guide* ]

---

## Layer 2 and Layer 3 Protocols

- **IPv4 for unicast forwarding**—In Junos OS Release 12.2, the ACX Series routers support basic IPv4 for unicast forwarding. The following key forwarding features are supported:
  - **Exception handling**—All basic exception handling features are supported, including but not limited to option packets, TTL expiry, MTU exceeded condition, redirect condition, and so on. In addition, Internet Control Message Protocol (ICMP) is supported to respond to various exception conditions.
  - **ARP**—Address Resolution Protocol (ARP) is supported to the full extent available in Junos OS, including but not limited to packet receive and transmit, ARP resolution trigger, and policing of ARP packets through implicit filters.
  - **IP fragmentation**—Fragmentation is in software and the number of packets fragmented is rate-limited.

[ See [TTL Processing on Incoming MPLS Packets](#) and [Configuring the Junos OS ARP Learning and Aging Options for Mapping IPv4 Network Addresses to MAC Addresses](#). ]

- **Layer 2 control packets**—The forwarding path supports the following types of Layer 2 control packets (excluding Operation, Administration, and Maintenance (OAM) packets) in both directions, receiving and forwarding:

- Ethernet control packets—ARP, IS-IS, 1588v2, Ethernet Synchronization Messaging Channel (ESMC).

[See [Configuring the Control Word for Layer 2 Circuits](#).]

- **Host path**—The host path to and from the CPU is supported in the following ways:
  - Host-bound traffic, prioritized into multiple queues, to support various levels of traffic.
  - Hardware-based policing used to limit denial-of-service attacks.
  - Protocol and flow-based policing.
  - Code point-based classification and prioritization of packets from the host to the external world.

[See [Path Messages](#).]

- **Keepalives**—The ACX Series routers support high resolution timers of up to 10 ms for driving keepalives for various OAM features, such as Bidirectional Forwarding Detection (BFD) and connectivity fault management (CFM).

[See [Junos OS Interfaces Fundamentals Configuration Guide](#).]

- **Counters and statistics**—Most packet-level and byte-level statistics for various entities in the forwarding path available in Junos OS are supported. The following counters and statistics are supported:
  - Ingress and egress packet and byte counters for logical interfaces, Ethernet pseudowires, and MPLS transit label-switched paths.
  - Discard packets counter for system-wide global Packet Forwarding Engine statistics.

[See [Display Traffic from the Point of View of the Packet Forwarding Engine](#).]

- **Statistics collection and reporting for Gigabit Ethernet interfaces**—For Gigabit Ethernet interfaces, Packet Forwarding Engine statistics are disabled by default. To enable Gigabit Ethernet interface statistics, you must specifically configure them. To configure Gigabit Ethernet interface statistics, include the **statistics** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level. To display statistics, issue the **show interfaces interface-name (brief | extensive)** operational mode command.

[See [Junos OS Ethernet Interfaces Configuration Guide](#) and [Fast Ethernet and Gigabit Ethernet Counters](#)]

- **Scaling and performance**—The following scaling and performance features are supported for interfaces and routes on the ACX Series routers:
  - **Interfaces**—Any logical interface enabled with IPv4 or MPLS is considered a Layer 3 interface. The maximum number of Layer 3 interfaces is 1000.
  - **Dual-tagged interfaces**—The Tag Protocol Identifier (TPID) for dual-tagged interfaces must meet the following conditions:
    - One inner TPID can be specified or used in the system.
    - The standard value of 0x8100 is allowed for the inner TPID.

- A maximum of four outer standard TPID values, that is, 0x8100, 0x9100, 0x9200, 0x88a8.
- **Route parameters**—On the ACX Series routers, all routes use a single, fully qualified match table and a single longest prefix match (LPM) route table. The following numbers assume an exclusive use of these tables for a particular type of route. If there is a mix, the numbers can change. The maximum number of supported routes is the following:
  - For IPv4, 8000 fully qualified match table and 12,000 LPM table.
  - For MPLS, 3000 label lookup entries, 2000 maximum transit unidirectional LSPs, and 1000 maximum Ethernet pseudowires. Only one MPLS lookup table is supported.



**NOTE:** Multicast is not supported on the ACX Series routers.



**NOTE:** With Junos OS, you can partition a single router into multiple logical devices that perform independent routing tasks. The ACX Series routers do not support this feature. Only one logical system is supported, the default logical system. The [edit logical-systems] hierarchy level is not supported.

- **Next-hop parameters**—The ACX Series router supports a maximum of 7000 unicast next-hop entries. This number is shared between IPv4, MPLS, and Ethernet pseudowires. The actual number is a little less than 7000 because a few of the next-hop entries are allocated and used internally. An additional 1000 of separate unicast entries are allowed for TDM and ATM pseudowires.
- **Address Resolution Protocol (ARP) parameters**—The maximum number of ARP entries is 7000.

[See [Junos OS Interfaces Fundamentals Configuration Guide](#).]

- **BFD and VCCV**—Bidirectional Forwarding Detection (BFD) support for virtual circuit connection verification (VCCV) enables you to configure a control channel for a pseudowire, in addition to the corresponding operations and management functions to be used over that control channel. BFD provides a low-resource mechanism for the continuous monitoring of the pseudowire data path and for detecting data plane failures.

[See [Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS](#).]



## MPLS

- **Label-switching router (LSR)**—With MPLS enabled, the ACX Series router can act as an LSR. An LSR processes label-switched packets and forwards packets based on their labels.

[See *Junos OS MPLS Applications Configuration Guide* and *MPLS Overview for ACX Series Universal Access Routers*.]

- **Label edge router (LER)**—The ACX Series router processes IPv4 traffic and pseudowire traffic over the MPLS network. The traffic is processed in both ingress and egress directions. Configuring MPLS on the LER is the same as configuring an LSR.

[See *Junos OS MPLS Applications Configuration Guide* and *MPLS Overview for ACX Series Universal Access Routers*.]

- **Pseudowire transport service**—A pseudowire carries Layer 1 and Layer 2 information over an IP/MPLS network infrastructure. Ethernet, ATM, and TDM pseudowires are supported. Only similar endpoints are supported on the ACX Series routers. For example, T1 to T1, ATM to ATM, and Ethernet to Ethernet.

[See *Pseudowire Overview for ACX Series Universal Access Routers*.]

- **Pseudowire redundancy**—A redundant pseudowire acts as a backup connection between PE routers and CE devices, maintaining Layer 2 circuits and services after certain types of failures. Pseudowire redundancy improves the reliability of certain types of networks (metro, for example) where a single point of failure could interrupt service for multiple customers. The following pseudowire redundancy features are supported:

- **Pseudowire standby**—A standby pseudowire can act as a backup connection between PE routers and CE devices, maintaining Layer 2 circuit and VPLS services after certain types of failures. To configure pseudowire standby, include the **backup-neighbor** statement at the [edit protocols l2circuit neighbor address interface *interface-name*] hierarchy level.

- **Protect interface**—A backup for the protected interface in case of failure. Network traffic uses the primary interface only so long as the primary interface functions. If the primary interface fails, traffic is switched to the protect interface. To configure the protect interface, specify the **protect-interface** statement at the [edit protocols l2circuit local-switching interface *interface-name*] hierarchy level.

- **Hot and cold standby**—Hot standby enables swift cutover to the backup or standby pseudowire. Cold standby is the inclusion of the **backup-neighbor** statement and the absence of the **standby** statement in the configuration. By default, a pseudowire is not backed up. The following hot standby configurations are supported:

- **Pseudowire hot standby**—A pseudowire configured with a backup neighbor is considered a standby pseudowire. When you configure that pseudowire with the **standby** statement at the [edit protocols l2circuit neighbor address interface *interface-name* **backup-neighbor**] hierarchy level, it is considered on hot standby. A pseudowire configured with only the **backup-neighbor** statement is considered on cold standby.

When you configure the **standby** statement on a backed-up pseudowire, traffic flows over both the active and standby pseudowires to the CE device. The CE device drops the traffic from the standby pseudowire, unless the active pseudowire fails. If the active pseudowire fails, the CE device automatically switches to the standby pseudowire.

- **Label-switched path (LSP) hot standby for secondary paths**—For an LSP, the hot standby state is meaningful only on secondary LSP paths. Maintaining a path in a hot-standby state enables swift cutover to the secondary path when downstream routers on the current active path indicate connectivity problems. To configure hot standby for an LSP, include the **standby** statement at the `[edit protocols mpls label-switched-path lsp-name secondary]` hierarchy level.
- **Ethernet connectivity fault management (CFM)**—The following major features of CFM for Ethernet pseudowires only are supported:
  - **Connection protection**—Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
  - **Path protection**—Path discovery and fault verification using the linktrace protocol. Similar to IP traceroute, this protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination.

[See *Redundant Pseudowires for Layer 2 Circuits and VPLS*, *Configuring the Protect Interface*, *Junos OS Layer 2 Configuration Guide*, and *Junos OS MPLS Applications Configuration Guide*.]

- **Control word**—The control word is 4 bytes long and is inserted between the Layer 2 protocol data unit (PDU) being transported and the virtual connection label. To configure the control word, include the **(control-word | no-control-word)** statement at the `[edit protocols l2circuit neighbor address interface interface-name]` hierarchy level.

[See *Configuring the Control Word for Layer 2 Circuits*.]

- **Uniform and pipe mode**—In an MPLS network, uniform mode is the default. Uniform mode makes all the nodes that a label-switched path (LSP) traverses visible to nodes outside the LSP tunnel. In contrast, pipe mode acts like a circuit and must be enabled. In pipe mode, when MPLS packets traverse the network, only the LSP ingress and egress points are visible to nodes that are outside the LSP tunnel. To configure pipe mode, include the **no-propagate-ttl** statement at the `[edit protocols mpls]` hierarchy level on each router that is in the path of the LSP. The global **no-propagate-ttl** statement disables time-to-live (TTL) propagation at the router level and affects all RSVP-signaled or LDP-signaled LSPs. Only the global configuration of TTL propagation is supported.

[See *no-propagate-ttl*.]

- **Exception packet handling for MPLS**—The following types of exception packet handling are supported:
  - Router alert
  - Time-to-live (TTL) expiry value
  - Virtual circuit connection verification (VCCV)

[See [Junos OS MPLS Applications Configuration Guide](#).]

- **Fast reroute**—Fast reroute is supported on ACX Series routers. Fast reroute provides redundancy for a label-switched path (LSP) path.

[See [Junos OS MPLS Applications Configuration Guide](#).]

- **Link protection**—Link protection helps ensure that traffic traversing a specific interface from one router to another can continue to reach its destination in the event that this interface fails.

[See [Link Protection](#).]

- **Node-link protection**—Node-link protection establishes a bypass LSP through a different router altogether.

[See [Node-Link Protection](#).]

- **MPLS ping and traceroute**—The ACX Series routers support MPLS ping and traceroute to the extent supported by Junos OS. Junos OS partially supports LSP ping and traceroute commands based on *RFC 4379, Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures*. However, Junos OS supports this functionality on LSP transit routers and head-end routers only. If a ping or traceroute command is issued from a router that fully supports RFC 4379, it can propagate correctly on routers running Junos OS.

[See [Pinging LSPs](#).]

---

## Network Management and Monitoring

- **Extends support for autoinstallation on ACX Series routers**—The autoinstallation mechanism for discovering, retrieving, and loading an appropriate configuration is now supported by the ACX Series Universal Access Routers.

[See [ACX Series Autoinstallation Overview](#).]

---

## Power Management

- **Power over Ethernet (PoE) (ACX2000 Universal Access routers)**—PoE is supported based on the IEEE 802.3af and IEEE 802.3at standards. Two ports on the ACX2000 router support PoE interfaces. The PoE interfaces permit electric power, along with data, to be passed over a copper Ethernet LAN cable. The PoE controller keeps track of the PoE power consumption on the router and allocates power to the PoE ports.
  - The PoE interface supports up to 65 W of Power over Ethernet Plus (PoE+).
  - With this new mode of power delivery, all four pairs of wires in the RJ45 cable have an option to deliver up to 65 W power per port provided high-power mode over the four pairs is requested. To enable high-power mode, include the **high-power** option at the `[edit poe management]` hierarchy level and include the **maximum-power watts** statement at the `[edit poe interface (interface-name | interface-all)]` hierarchy level.
  - Control the PoE interfaces with the following configuration statements and commands:

- To enable PoE physical interfaces, include the **interface** statement at the **[edit poe]** hierarchy level. Specify an individual PoE interface with the **interface-name** option, or all PoE interfaces with the **interface-all** option.
- Disable the PoE interface with the **disable** statement at the **[edit poe interface-name | interface-all]** hierarchy level.
- Configure the PoE interface to gather voltage and power information by including the **telemetries** statement at the **[edit poe interface (interface-name | interface-all)]** hierarchy level. Specify the following options for this statement: **disable**, **duration hours**, and **interval minutes**.
- Display the power consumption with the **show poe controller** command.
- Display the configured PoE interfaces with the **show poe interface** command.

[See *Understanding PoE on ACX Series Universal Access Routers*, *Junos OS Ethernet Interfaces Configuration Guide*, and *Junos OS System Basics Configuration Guide*.]

---

## Routing Policy and Firewall Filters

- **Firewall features supported on ACX Series Universal Access Routers**—Existing Junos OS firewall features are supported without changes to statements or functionality.

The following is the list of key supported firewall features and any conditions associated with them:

- Configuration of filters for the following protocol families only: **any**, **ccc**, **inet**, and **mpls**.
- Firewall filters applied to a logical interface must have the **interface-specific** statement included at the respective family hierarchy level.
- An egress filter must always have the **interface-specific** statement configured.
- Configuration of policers and three-color policers.
- Actions—for example, **count**, **discard**, **log**, and so on.
- Operational mode commands for firewall filters are supported on the ACX Series routers without changes.

[See *Standard Firewall Filter Match Conditions and Actions on ACX Series Routers Overview*.]

- **Filter-based forwarding for routing instances**—For IPv4 traffic only, you can use stateless firewall filters in routing instances to control how packets travel in a network. This is called filter-based forwarding.

You can define a firewall filtering term that directs matching packets to a specified routing instance. This type of filtering can be configured to route specific types of traffic through a firewall or other security device before the traffic continues on its path. To configure a stateless firewall filter to direct traffic to a routing instance, configure a term with the **routing-instance routing-instance-name** terminating action at the **[edit firewall family inet filter filter-name term term-name then]** hierarchy level to specify the routing instance to which matching packets will be forwarded. To configure the filter to direct traffic to the master routing instance, use the **routing-instance default**

statement at the **[edit firewall family inet filter *filter-name* term *term-name* then]** hierarchy level.

*[ACX Series Universal Access Router Configuration Guide]*

- **Forwarding table filters for routing instances**—Forwarding table filter is a mechanism by which all the packets forwarded by a certain forwarding table are subjected to filtering and if a packet matches the filter condition, the configured action is applied on the packet. You can use the forwarding table filter mechanism to apply a filter on all interfaces associated with a single routing instance with a simple configuration. You can apply a forwarding table filter to a routing instance of type forwarding and also to the default routing instance **inet.0**. To configure a forwarding table filter, include the **filter *filter-name*** statement at the **[edit firewall family inet]** hierarchy level.

*[ACX Series Universal Access Router Configuration Guide]*

## Routing Protocols

- **Support for Layer 3 VPNs for IPv4 and IPv6 address families**—You can configure Layer 3 virtual private network (VPN) routing instances on ACX Series routers at the **[edit routing-instances *routing-instance-name* protocols]** hierarchy level for unicast IPv4, multicast IPv4, unicast IPv6, and multicast IPv6 address families. If you do not explicitly specify the address family in an IPv4 or an IPv6 environment, the router is configured to exchange unicast IPv4 or unicast IPv6 addresses by default. You can also configure the router to exchange unicast IPv4 and unicast IPv6 routes in a specified virtual routing and forwarding (VRF) routing instance. If you specify the multicast IPv4 or multicast IPv6 address family in the configuration, you can use BGP to exchange routing information about how packets reach a multicast source, instead of a unicast destination, for transmission to endpoints.

Only the forwarding and virtual router routing instances support unicast IPv6 and multicast IPv6 address families. Unicast IPv6 and multicast IPv6 address families are not supported for VRF routing instances.

A VRF routing instance is a BGP and MPLS VPN environment in which BGP is used to exchange IP VPN routes and discover the remote site, and VPN traffic traverses an MPLS tunnel in an IP and MPLS backbone. You can enable an ACX Series router to function as a provider edge (PE) router by configuring VRF routing instances.

You can configure the following types of Layer 3 routing instances:

- **Forwarding**—Use this routing instance type for filter-based forwarding applications.
- **Virtual router**—A virtual router routing instance is similar to a VRF instance type, but is used for non-VPN-related applications.
- **VRF**—Use the VRF routing instance type for Layer 3 VPN implementations. This routing instance type has a VPN routing table as well as a corresponding VPN forwarding table. For this instance type, there is a one-to-one mapping between an interface and a routing instance. Each VRF routing instance corresponds with a forwarding table. Routes on an interface go into the corresponding forwarding table. This routing instance type is used to implement BGP or MPLS VPNs in service provider networks or in big enterprise topologies.

[*ACX Series Universal Access Router Configuration Guide*]

- **Support for Multiprotocol BGP**—Multiprotocol BGP (MBGP) is an extension to BGP that enables BGP to carry routing information for multiple network layers and address families. MBGP can carry the unicast routes used for multicast routing separately from the routes used for unicast IP forwarding.

You can configure MBGP on ACX Series routers for IPv4 and IPv6 address families in the following ways:

- To enable MBGP to carry network layer reachability information (NLRI) for address families other than unicast IPv4, include the **family inet** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MBGP to carry NLRI for the IPv6 address family, include the **family inet6** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MBGP to carry Layer 3 virtual private network (VPN) NLRI for the IPv4 address family, include the **family inet-vpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MBGP to carry Layer 3 VPN NLRI for the IPv6 address family, include the **family inet6-vpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MBGP to carry multicast VPN NLRI for the IPv4 address family and to enable VPN signaling, include the **family inet-mvpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.
- To enable MBGP to carry multicast VPN NLRI for the IPv6 address family and to enable VPN signaling, include the **family inet6-mvpn** statement at the **[edit protocols bgp]** or the **[edit routing-instances routing-instance-name protocols bgp]** hierarchy level.

[*ACX Series Universal Access Router Configuration Guide*]

---

## Software Architecture

- **ACX Series router architecture**—The ACX Series router is a single-board router with a built-in Routing Engine and one Packet Forwarding Engine that has one Flexible PIC Concentrator (FPC 0). Because there is no switching fabric, the single Packet Forwarding Engine takes care of packet forwarding.
  - Routing Engine—Provides Layer 3 routing services and network management.
  - Packet Forwarding Engine—Performs Layer 2 and Layer 3 packet switching, route lookups, and packet forwarding.

[See [ACX Series Universal Access Router Overview](#).]
- **Packet Forwarding Engine management**—The **request chassis feb restart slot slot-number** command is introduced to restart the specified Forwarding Engine Board

(FEB). When you enter this command, you are provided feedback on the status of your request. For example:

```
user@host> request chassis feb restart slot 0
FEB will be restarted NOW.
```

[See [request chassis feb](#).]

- **Dual-speed Gigabit Ethernet interface**—The Gigabit Ethernet ports on the router have the capacity to work as a 1- or 10-Gigabit Ethernet interface, depending on the type of small form-factor pluggable (SFP) transceiver inserted. When you insert an SFP+ transceiver, the interface works at the 10-gigabit speed. When you insert an SFP transceiver, the interface works at the 1-gigabit speed. Configuration is not required because the speed is determined automatically based on the type of inserted SFP transceiver. The dual-speed interface is automatically created with the **xe** prefix, for example, **xe-4/0/0**.

The same configuration statements are used for both speeds, and CoS parameters are scaled as a percentage of the port speed. To configure a dual-speed Gigabit Ethernet interface, include the **interface xe-fpc/pic/port** statement at the [edit interfaces] hierarchy level. To display the interface speed and other details, issue the **show interfaces** command.

[See [Understanding Interfaces on ACX Series Universal Access Routers](#).]

- **SNMP and MIB support**—The ACX Series routers support all existing MIBs that identify all the different components of the chassis—for instance, the power supply. Existing MIB support is defined in [Standard SNMP MIBs Supported by Junos OS](#) and [Enterprise-Specific MIBs and Supported Devices](#).
- **Memory utilization**—The **show chassis routing-engine** and the **show chassis feb** commands can be used to find the memory allocated for each of the Routing Engine and Packet Forwarding Engine components.

[See [show chassis routing-engine](#) and [show chassis feb](#).]

- **System snapshot support**—The **request system snapshot** command enables you to create a copy of the currently running software on another media—for example, a universal serial bus (USB) storage device, the active slice of a dual-root partitioned router, or the alternate slice of a dual-root partitioned router. Typically, this command is used prior to the upgrade of the software image on the dual internal NAND flash device (with the **da0s1** or **da0s2** slices) or to remedy a bad image, thereby preventing the bad image from rendering the system useless. A snapshot to another media ensures that the device can boot from the other media in case the system does not boot from the current image.

[See [Understanding System Snapshot on an ACX Series Router, Example: Taking a Snapshot of the Software and Configuration](#), and [request system snapshot \(ACX Series\)](#).]

## Timing and Synchronization

- **Timing and synchronization support at the chassis level**—All existing Junos OS timing and synchronization features are supported at the [edit chassis synchronization] hierarchy level without changes to statements or functionality, except for the **external-a**

and the **external-b** statements, which are not supported on the ACX Series routers. Instead of the **external-a** and the **external-b** statements, the ACX Series routers support the new **bits** and **gps** statements at the **[edit chassis synchronization source]** hierarchy level.

- **bits**—The external building-integrated timing supply (BITS) device is connected to the router's T1 or E1 BITS interface, which upon configuration becomes a candidate for selection as the clock source by the clock source selection algorithm.
- **gps**—The 10-MHz clock input received from the Global Positioning System (GPS) is considered one of the candidate sources for chassis synchronization by the clock source selection algorithm.

Both the **bits** and **gps** statements include the following options:

- **priority number**—Specify a priority level between 1 and 5. When not specified, **gps** has a higher default priority than **bits**, and **bits** has a higher default priority than other Gigabit Ethernet, 10-Gigabit Ethernet, T1, or E1 clock sources, which have the lowest default priority.
- **quality-level (prc | prs | sec | smc | ssu-a | ssu-b | st2 | st3 | st3e | st4 | stu | tnc)**—Specify the expected quality of the incoming clock on this source. Specific **quality-level** options are valid depending on the configured **network-option**: **option-1** or **option-2** at the **[edit chassis synchronization]** hierarchy level.
  - Both option I and option II SSM quality levels (QL) are supported:
    - Both **option-1** and **option-2** Synchronization Status Message (SSM) quality levels (QL) are supported:
    - For option-2, the default QL for external clocks is QL\_STU whether or not QL is enabled.
- **request force-switch**—Force a switch to the source provided that the source is enabled and not locked out. Only one configured source can be force-switched.
- **request lockout**—A lockout can be configured for any source. When a lockout is configured for a source, that source will not be considered by the selection process.

[See [Clock Sources for the ACX Series Universal Access Routers](#), *bits*, and *gps*.]

- **T1 or E1 BITS interface (ACX2000 router)**—The ACX2000 router has a T1 or E1 building integrated timing source (BITS) interface that you can connect to an external clock. After you connect the interface to the external clock, you can configure the BITS interface so that the BITS interface becomes a candidate source for chassis synchronization to the external clock. The frequency of the BITS interface depends on the Synchronous Ethernet equipment (EEC) slave clock selected with the **network-option** statement at the **[edit chassis synchronization]** hierarchy level.
  - **option-1**—EEC—Option 1 applies to Synchronous Ethernet equipment optimized for 2048 Kbps. With this option, the BITS interface operates at the speed of an E1 interface.



- **option-2**—EEC-Option 2 applies to Synchronous Ethernet equipment optimized for 1544 Kbps. With this option, the BITS interface operates at the speed of a T1 interface.

To configure the BITS interface as the candidate source for synchronization, include the **bits** statement and options at the **[edit chassis synchronization source]** hierarchy level.

[See [External Clock Synchronization Overview for ACX Series Routers](#) and [source \(Chassis Synchronization\)](#).]

- **Global Positioning System (GPS)**—GPS is a navigation aid system that uses signals from satellites to calculate the actual position of a GPS-capable receiver. These signals are not only used for determining the position of the receiver on Earth but also as a very accurate time base. There are GPS receivers with 10-MHz clock frequency output synchronized to a GPS satellite. The ACX Series router has a SubMiniature version B (SMB) connector that can take 10-MHz sine-wave input from a GPS receiver. To configure this 10-MHz clock from a GPS receiver as a candidate clock source for chassis synchronization, include the **gps** statement and options at the **[edit chassis synchronization source]** hierarchy level.

[See [Configuring External Clock Synchronization for ACX Series Routers](#) and [gps](#).]

- **Automatic clock selection**—In automatic clock selection, the system chooses up to two best upstream clock sources. The system then uses the clock recovered from one of the sources to lock the chassis clock. If an upstream clock with acceptable good quality is not available or if the system is configured in free-run mode, the system uses the internal oscillator. The following automatic clock selection features are supported for Synchronous Ethernet, T1 or E1 line timing sources, and external inputs:



**NOTE:** Automatic clock selection does not apply to the IEEE 1588v2 recovered clock.

- **Basis of automatic clock selection**—Automatic clock selection of the best quality clock source is based on the Ethernet Synchronization Message Channel (ESMC) Synchronization Status Message (SSM) quality level, the configured quality level, and the priority. To configure the clock mode, include the **clock-mode** statement with the **free-run** option or the **auto-select** option at the **[edit chassis synchronization]** hierarchy level. When the **free-run** option is configured, the chassis is locked to the free-running local oscillator, which is the Stratum 3E oscillator. The **auto-select** option enables the clock source selection algorithm to run.

[See [clock-mode](#).]

- **Clock source selection algorithm**—The clock source selection algorithm is triggered by the following events:
  - Signal failure detected on the currently selected source.
  - Changes in the received ESMC SSM quality level.
  - Configuration changes. For example, the addition or deletion of a clock source, a change to the quality level mode, and so on.

Automatic clock selection supports two modes on the ACX Series router: QL enabled and QL disabled. To configure QL mode, include the **quality-mode-enable** statement at the **[edit chassis synchronization]** hierarchy level.

- **QL disabled**—The default setting is disable, which means that when the **quality-mode-enable** statement is not configured, quality level is disabled. In this mode, the best clock is selected based on the configured ESMC SSM quality level. If the quality levels of the configured clocks are equal, the clock selection is based on the configured priority. If both the configured quality levels and priority are equal, one of the sources is randomly selected.
- **QL enabled**—In this mode, the best clock is selected based on the incoming ESMC SSM quality level as long as the incoming quality level is at least as good as the source's configured quality level. If the quality levels are equal, the clock selection is based on the configured priority. If both the received quality level and the priority are equal, one of the sources is selected randomly.
- **Configured or received clock selection**—The **selection-mode (configured-quality | received-quality)** statement specifies whether the clock source selection algorithm should use the configured or received ESMC SSM quality level for clock selection. In both the selection modes, the interface qualifies for clock source selection only when the received ESMC SSM quality level on the interface is equal to or greater than the configured ESMC SSM quality level for the interface.

When the **selection-mode** statement is set as **configured-quality**, the clock source selection algorithm uses the ESMC SSM quality level configured for a clock source.

When the **selection-mode** statement is set as **received-quality**, the clock source selection algorithm uses the ESMC SSM quality level received on the interface that is configured as a clock source.



**NOTE:** For the **selection-mode** statement configuration to take effect, you must set the **quality-mode-enable** statement at the **[edit chassis synchronization]** hierarchy level.

---

[See [Automatic Clock Selection Overview](#), [Clock Sources for the ACX Series Universal Access Routers](#), and [synchronization \(ACX Series\)](#).]

- **Synchronous Ethernet (ACX2000 router)**—Synchronous Ethernet is a physical layer frequency transfer technology modeled after synchronization in SONET/SDH. Traditional Ethernet nodes, which do not support Synchronous Ethernet, do not carry synchronization from one node link to another. Synchronous Ethernet capable nodes, however, can synchronize their chassis clock to a clock recovered from an interface connected to an upstream clock master. After which, the clock is used to time data sent to downstream clock slaves, forming a synchronization trail from a primary reference clock (PRC) to Ethernet equipment clocks (EECs) and transferring frequency synchronization along the trail.

The ITU G.8264 specification defines the Synchronization Status Message (SSM) protocol and its format for Synchronous Ethernet to ensure interoperability between Synchronous Ethernet equipment used for frequency transfer—for example,

SONET/SDH. Synchronous Ethernet provides stable frequency synchronization to a PRC and is not affected by load on the network. However, it requires that all the nodes from the PRC to the last downstream node are Synchronous Ethernet capable. Synchronous Ethernet is a recommended technology for mobile networks that require frequency-only synchronization—for example, 2G or 3G base stations.

[See [Synchronous Ethernet Overview on the ACX Series Universal Access Routers](#).]

- **Precision Timing Protocol (PTP), also known as IEEE 1588v2**—PTP synchronizes clocks between nodes in a network, thereby enabling the distribution of an accurate clock over a packet-switched network. This synchronization is achieved through packets that are transmitted and received in a session between a master clock and a slave clock. The master clock is external to the ACX Series router, for example, a TCA Series Timing Client or an MX Series router.

Most existing PTP statements are supported without changes in functionality, see *[edit protocols ptp] Hierarchy Level* for details about particular statements. The following new PTP statements are supported:

- **ipv4-dscp number**—Specifies the value used as the DiffServ code point (DSCP) value for all PTP IPv4 packets originated by the router. To configure the DSCP value, include the **ipv4-dscp number** statement at the *[edit protocols ptp]* hierarchy level.

[See [ipv4-dscp](#).]

- **announce-interval announce-interval-value**—This value specifies the rate of announce messages that a PTP slave clock requests from the master clock during a unicast negotiation session. The announce interval is configured on the slave clock. To configure the announce interval, include the **announce-interval announce-interval-value** statement at the *[edit protocols ptp slave]* hierarchy level. The configuration of the **announce-interval** statement is effective only when the **unicast-negotiation** statement is also configured at the *[edit protocols ptp]* hierarchy level.

[See [announce-interval](#).]

- **grant-duration interval**—When unicast negotiation is enabled, the local PTP slave clock requests announce, sync, and delay-response messages from the master clock. In each request, the slave clock asks for the packets to be sent at a specified rate, and it provides a duration for which the rate is valid. The grant-duration value is specified in seconds. The default grant duration is 3600 seconds or 1 hour. To configure the grant duration, include the **grant-duration interval** statement at the *[edit protocols ptp slave]* hierarchy level.

[See [grant-duration](#).]

- **asymmetry number**—A compensating value for networks in which there is path asymmetry between the 1588v2 slave and master clocks. Specify a positive or negative value that is added to the path delay value from the slave clock to the master clock, making the delay symmetric and equal to the path from the master clock to the slave clock. The asymmetry value is in nanoseconds and can vary from minus (–) 100 milliseconds to 100 milliseconds, allowing compensation for up to 1/10 of a second of path asymmetry. To configure an asymmetrical value, include the **asymmetry number** statement at the *[edit protocols ptp slave interface]*

*interface-name* unicast-mode clock-source *ip-address* local-ip-address *ip-address*] hierarchy level.

[See [asymmetry](#).]

- **sync-interval *interval***—Requested log mean interval between sync messages. The **sync-interval** is configured on the slave clock and specifies the rate at which sync messages are requested to be sent from the master clock to the slave clock. The specified value is the log2 value of the requested sync packet rate. Because the accepted value varies from -6 to 0, the specified packet rate will be from  $2^{-6}$  to  $2^0$  or from 64 packets per second to 1 packet per second.

The configuration of the **sync-interval** statement is effective only when the **unicast-negotiation** statement is also configured at the [edit protocols ptp] hierarchy level.

[See [sync-interval](#).]

The following key PTP features are supported:

- **Ordinary clock (slave only)**—The PTP ordinary slave clock estimates time offset from the PTP master clock and tries to align its own time and frequency with that of the master clock. ACX Series routers support the IEEE 1588v2 compliant ordinary slave clock. To configure a slave clock, include the **slave** statement and options at the [edit protocols ptp] hierarchy level.
- **PTP over UDP over IPv4**—The IEEE1588v2 standard specifies different transport protocols for carrying PTP packets. For example, PTP over Ethernet, PTP over UDP over IPv4, and PTP over UDP over IPv6. The ACX Series routers support PTP over UDP over IPv4.
- **Unicast mode (IPv4 on Gigabit Ethernet interfaces only)**—Unicast mode is a user-to-user protocol used to send a datagram to a single recipient. Unicast mode is used for transporting PTP messages. To configure unicast mode on an interface, include the **unicast-mode** statement at the [edit protocols ptp slave interface *interface-name*] hierarchy level.

[See [Precision Timing Protocol \(PTP\) on ACX Series Universal Access Routers](#), [edit protocols ptp] Hierarchy Level, [Example: Configuring an Ordinary Slave Clock With Unicast-Negotiation](#), and [Example: Configuring an Ordinary Slave Clock Without Unicast-Negotiation](#).]

#### Related Documentation

- [Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers on page 34](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Resolved Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

## Changes in Default Behavior and Syntax in Junos OS Release 12.2 for ACX Series Routers

- [IPv6 on page 29](#)

### IPv6

- **Change in automatically generated virtual-link-local-address for VRRP over IPv6—** Over IPv6 will be 0x02. This change makes the VRRP over IPv6 feature in Junos OS 12.2R5, 12.3R3, 13.1R3, and later releases inoperable with Junos OS 12.2R1, 12.2 R2, 12.2 R3, 12.2R4, 12.3R1, 12.3R2, 13.1R1, and 13.3R2 releases if the automatically generated virtual-link-local-address ID used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

#### Related Documentation

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

## Known Limitations in Junos OS Release 12.2 for ACX Series Routers

The following software limitations currently exist in Juniper Networks ACX Series Universal Access Routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

### Class of Service (CoS)

- When the **rewrite-rules** statement is configured with the **dscp** or the **inet-precedence** option at the **[edit class-of-service interfaces]** hierarchy level, the expectation is that the DiffServ code point (DSCP) or IPv4 precedence rewrite rules take effect only on IP packets. However, in addition to the IP packets, the DSCP or IPv4 rewrite takes effect on the IP header inside the Ethernet pseudowire payload as well. [[PR664062](#): This is a known limitation.]

### Interfaces and Chassis

- When the **differential-delay *number*** option is configured in the **ima-group-option** statement at the **[edit interfaces at-fpc/pic/ima-group-no]** hierarchy level, with a value less than 10, some of the member links might not come up and the group might remain down resulting in traffic loss. A workaround is to keep the differential delay value above 10 for all IMA bundles. [[PR726279](#): This is a known limitation.]
- On ACX1000 and ACX2000 routers, the outbound host traffic does not show the respective interface queue statistics as per the configuration. This is due to a hardware limitation. However, the actual queuing and scheduling happens as per the configuration. [PR772149](#): This is a known limitation.]

### Multiprotocol Label Switching (MPLS)

---

- The scaling numbers for pseudowires and MPLS label routes published for the ACX Series routers are valid only when the protocols adopt graceful restart. In case of non-graceful restart, the scaling numbers would become half of the published numbers. [PR683581: This is a known limitation.]

### Routing Policy and Firewall Filters

---

- On ACX Series routers, packet drops in the egress interface queue are also counted as *input packet rejects* under the **Filter statistics** section in the output of the **show interface extensive** command when it is run on the ingress interface. [PR612441: This is a known limitation.]
- When the **statistics** statement is configured on a logical interface—for example, the [edit interface name-X unit unit-Y]; when the (**policer** | **count** | **three-color-policer**) statements are configured in a firewall filter for the **family any**—for example, the [edit firewall family any filter filter-XYZ term term-T then] hierarchy level; and when the configured **filter-XYZ** is specified in the **output** statement of the logical interface at the [edit interface name-X unit unit-Y filter] hierarchy level, the counters from the configuration of another firewall family filter on the logical interface do not work. [PR678847: This is a known limitation.]
- The policing rate can be incorrect if the following configurations are applied together:
  - The **policer** or **three-color-policer** statement configured in a firewall filter—for example, **filter-XYZ** at the [edit firewall family any filter filter-XYZ term term-T then] hierarchy level, and **filter-XYZ** is specified as an ingress or egress firewall filter on a logical interface—for example, **interface-X unit-Y** at the [edit interface interface-X unit unit-Y filter (input|output) filter-XYZ] hierarchy level.
  - The **policer** or **three-color-policer** statement configured in a firewall filter—for example, **filter-ABC** at the [edit firewall family name-XX filter filter-ABC term term-T then] hierarchy level, and **filter-ABC** is configured as an ingress or egress firewall filter on a family of the same logical interface **interface-X unit-Y** at the [edit interface interface-X unit unit-Y family name-XX filter (input|output) filter-ABC] hierarchy level.



**NOTE:** If one of these configurations is applied independently, then the correct policer rate is observed.

---

[PR678950: This is a known limitation.]

#### Related Documentation

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers on page 34](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Resolved Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)

- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

## Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers

The following problems currently exist in Juniper Networks ACX Series Universal Access Routers. The identifier following the descriptions is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

### Interfaces and Chassis

---

- On ACX1000 and ACX2000 routing platforms outbound host traffic does not show the respective interface queue statistics as per the configuration. This is due to a hardware limitation. However, the actual queuing and scheduling happens as per the configuration. [PR772149](#)
- ACX does not support IFL statistics for IFLs with vlan-list/vlan-range. [PR810973](#)

### Related Documentation

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers on page 34](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Resolved Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

## Resolved Issues in Junos OS Release 12.2 for ACX Series Routers

### Current Release

---

The following are the issues that have been resolved in Junos OS Release 12.2R9 for Juniper Networks ACX Series Universal Access Routers. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

- [Release 12.2R9 on page 31](#)

### **Release 12.2R9**

- Note: There are no resolved issues in Junos OS Release 12.2R9.

### Previous Releases

---

The following are the issues that have been resolved since Junos OS Release 12.2R1. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system:

- [Release 12.2R8 on page 32](#)
- [Release 12.2R7 on page 32](#)

- [Release 12.2R6 on page 32](#)
- [Release 12.2R5 on page 32](#)
- [Release 12.2R4 on page 32](#)
- [Release 12.2R3 on page 32](#)
- [Release 12.2R2 on page 33](#)
- [Release 12.2R1 on page 33](#)

### **Release 12.2R8**

#### ***Subscriber Access Management***

- jdhcpd crash triggered by DHCP client re-discovering on different vlan. [ [PR913823](#): This issue has been resolved.]

### **Release 12.2R7**

#### ***Interfaces and Chassis***

- Heater status shows to be on under normal operating temperatures. This is a cosmetic issue and has been addressed. [ [PR898079](#): This issue has been resolved.]

### **Release 12.2R6**

#### ***MPLS***

- A commit for a configuration change that simultaneously disables RSVP and a point-to-point interface (like so, t1, atm) might generate an rpd core file. To solve this issue, do not commit a configuration change that simultaneously disables RSVP and a point-to-point interface. Rather disable RSVP and point-to-point interfaces in separate configuration commits. [ [PR782174](#): This issue has been resolved.]

### **Release 12.2R5**

#### ***MPLS***

- This is just an SNMP Trap generated due to inconsistency in temperature parameters and thresholds. [ [PR859507](#): This issue has been resolved.]

### **Release 12.2R4**

#### ***VLAN Infrastructure***

- On a Layer 2 circuit, when any one of the logical interfaces on the PE routers is disabled after changing VLAN-tagging to flexible VLAN tagging or vice versa on a CE router, the pseudowires return an MTU mismatch error. As a workaround, disable and then enable all the logical interfaces on the PE routers. [ [PR834466](#): This issue has been resolved.]

### **Release 12.2R3**

#### ***Interfaces and Chassis***

- Clock source when configured on 0th port on the Gigabit Ethernet interface (as mentioned below) would result in link flaps. The rest of the ports, when configured with clock source, work fine.



Avoid using clock source configuration on the following ports on the following ACX variants:

- ACX1000: ge-0/2/0
- ACX1100: ge-0/1/0
- ACX2000: ge-0/2/0
- ACX2100: ge-1/0/0, ge-1/1/0, and ge-1/2/0
- ACX2200: ge-0/0/0, ge-0/1/0, and ge-0/2/0

[[PR827449](#): This issue has been resolved.]

### **Release 12.2R2**

#### **User Interface and Configuration**

- In Junos OS Release 12.2R2 and later, the following are applicable:
  - CLI restriction for configuring valid value of data-tlv-size for SLA iterator based delay measurement (which is available): **set protocols oam ethernet connectivity-fault-management maintenance-domain <md-name> maintenance-association <ma-name> mep <mep-id> remote-mep <remote-mep-id> sla-iterator-profile <sla-profile-name> data-tlv-size <value>**  
The valid range is changed from 1...1400 to 64...1400.
  - CLI restriction for configuring valid value of data in data TLV of delay measurement request packets (which is available): **run monitor ethernet delay-measurement two-way size <value>**  
The valid range is changed from 1...1400 to 64...1400.

[[PR800605](#): This issue has been resolved.]

### **Release 12.2R1**

#### **Interfaces and Chassis**

- When you issue the **show system memory** command on MX80 routers, the "**unable to load pmap\_helper module: No such file or directory**" error message is displayed in the output of the command. [[PR737616](#): This issue has been resolved.]
- Synchronous Ethernet on ACX Series routers might not work when the selected source port is other than ge-0/1/0. [[PR751695](#): This issue has been resolved.]
- An FPC restart affecting point-to-point interfaces (such as so, tl, or atm) might generate a routing protocol process (rpd) core file. As a workaround, do not restart an FPC with one of these interfaces. [[PR782174](#): This issue has been resolved.]

### MPLS

- When an ACX Series Universal Access Router is working as a transit router in an MPLS network, then traceroute response from the router displays "nexthop" information as "Unhelpful". [[PR669005](#): This issue has been resolved.]

### Related Documentation

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers on page 34](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

## Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers

- [Hardware](#)
- [Class of Service \(CoS\)](#)
- [Routing Protocols](#)

### Hardware

---

- In the "ACX2000 and ACX2100 DC Power Specifications" topic of the *ACX2000 and ACX2100 Router Hardware Guide*, the DC input voltages row in the table presented in the topic incorrectly mentions that the range is 18 to 30 VDC. The correct DC input voltage range for a nominal 24-volt operation is 20 to 30 VDC.

### Class of Service (CoS)

---

- Support for multifold classifiers is incorrectly omitted from the ACX documentation. Multifold classifiers allow fine grained classification by examination of multiple fields in the packet header—for example, the source and destination address of the packet, and the source and destination port numbers of the packet. A multifold classifier typically matches one or more of the six packet header fields: destination address, source address, IP protocol, source port, destination port, and DSCP. Multifold classifiers are used when a simple BA classifier is insufficient to classify a packet.

In the Juniper Networks Junos operating system (Junos OS), you configure a multifold classifier with a firewall filter and its associated match conditions. This enables you to use any filter match criteria to locate packets that require classification. From a CoS perspective, multifold classifiers (or firewall filter rules) provide the following services:

- Classify packets to a forwarding class and loss priority. The forwarding class determines the output queue. The loss priority is used by schedulers in conjunction with the random early discard (RED) algorithm to control packet discard during periods of congestion.

- Police traffic to a specific bandwidth and burst size. Packets exceeding the policer limits can be discarded, or can be assigned to a different forwarding class, to a different loss priority, or to both.



**NOTE:** You police traffic on input to conform to established CoS parameters, setting loss handling and forwarding class assignments as needed. You shape traffic on output to make sure that router resources, especially bandwidth, are distributed fairly. However, input policing and output shaping are two different CoS processes, each with their own configuration statements.

To configure multifield classifiers, include the following statements at the `[edit firewall]` hierarchy level:

```
[edit firewall]
family family-name {
  filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        dscp 0;
        forwarding-class class-name;
        loss-priority (high | low);
      }
    }
  }
  simple-filter filter-name {
    term term-name {
      from {
        match-conditions;
      }
      then {
        forwarding-class class-name;
        loss-priority (high | low | medium);
      }
    }
  }
}
```

The minimum configuration required to define a multifield classifier is the following:

```
[edit firewall]
family family-name {
  simple-filter filter-name {
    term term-name {
      then {
        forwarding-class class-name;
        loss-priority (high | low | medium);
      }
    }
  }
}
```

After defining the multifield classifier, you can apply the multifield classifier to an individual interface with the following configuration:

```
[edit interfaces]
interface-name{
  unit logical-unit-number{
    family family {
      filter {
        input filter-name;
      }
    }
  }
}
```

[ACX Series Universal Access Router Configuration Guide]

---

### Routing Protocols

- The following additional information about the support for unicast IPv6 and multicast IPv6 address families for routing instances on ACX Series routers applies to the *Routing Protocols* subsection in the *New Features in Junos OS Release 12.2 for ACX Series Routers* section of the Junos OS 12.2R2 Release Notes:

Only the forwarding and virtual router routing instances support unicast IPv6 and multicast IPv6 address families. Unicast IPv6 and multicast IPv6 address families are not supported for VRF routing instances.

[*Release Notes*]

#### Related Documentation

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Resolved Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers on page 36](#)

## Upgrade and Downgrade Instructions for Junos OS Release 12.2 for ACX Series Routers

This section discusses the following topics:

- [Basic Procedure for Upgrading to Release 12.2 on page 36](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 39](#)

---

### Basic Procedure for Upgrading to Release 12.2

When upgrading or downgrading Junos OS, 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:** Before upgrading, back up the file system and the currently active Junos OS 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 Junos OS. 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 *Understanding System Snapshot on an ACX Series Router*.

.....

The download and installation process for Junos OS Release 12.2 is different from previous Junos OS releases. Follow these steps:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks web page:  
<http://www.juniper.net/support/downloads/>
2. Select the name of the Junos platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. 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-12.2R9-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-12.2R9-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.

Including 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 OS Release 12.2 `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 nonemergency 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 *Junos OS Multiplay Solutions Guide*.

---

### Upgrade and Downgrade Support Policy for Junos OS Releases

---

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

**Related  
Documentation**

- [New Features in Junos OS Release 12.2 for ACX Series Routers on page 5](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for ACX Series Routers on page 34](#)
- [Known Limitations in Junos OS Release 12.2 for ACX Series Routers on page 29](#)
- [Outstanding Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)
- [Resolved Issues in Junos OS Release 12.2 for ACX Series Routers on page 31](#)



## Junos OS Release Notes for EX Series Switches

---

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches on page 49](#)
- [Limitations in Junos OS Release 12.2 for EX Series Switches on page 50](#)
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches on page 57](#)
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches on page 64](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches on page 92](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches on page 94](#)

### New Features in Junos OS Release 12.2 for EX Series Switches

This section describes new features in Release 12.2 of the Junos operating system (Junos OS) for EX Series switches.

Not all EX Series software features are supported on all EX Series switches in the current release. For a list of all EX Series software features and their platform support, see [EX Series Switch Software Features Overview](#) and [EX Series Virtual Chassis Software Features Overview](#).

New features are described on the following pages:

- [Hardware on page 42](#)
- [Access Control and Port Security on page 43](#)
- [Ethernet Switching and Spanning Trees on page 44](#)
- [Firewall Filters on page 45](#)
- [High Availability on page 45](#)
- [Infrastructure on page 46](#)
- [Interfaces on page 46](#)
- [J-Web Interface on page 47](#)
- [Management and RMON on page 47](#)
- [MPLS on page 47](#)
- [Power over Ethernet \(PoE\) on page 48](#)
- **Software Installation and Upgrade**
- [Virtual Chassis on page 48](#)

## Hardware

---

- **EX2200 Virtual Chassis**—You can now configure an EX2200 Virtual Chassis. A Virtual Chassis is a collection of interconnected switches that you can manage and control as a single entity through one master switch. You can interconnect up to four EX2200 switches in a Virtual Chassis. [See [EX2200](#), [EX3300](#), [EX4200](#), [EX4500](#), and [EX4550 Virtual Chassis Overview](#).]
- **EX4550 switches**—EX4550 switches provide high performance, scalable connectivity, and carrier-class reliability for high-density environments such as campus aggregation, branch offices, and data center networks.

EX4550 switches are available in the following eight models, with AC or DC power supplies:

- EX4550-32F-AFI
- EX4550-32F-AFO
- EX4550-32F-DC-AFI
- EX4550-32F-DC-AFO
- EX4550-32T-AFI
- EX4550-32T-AFO
- EX4550-32T-DC-AFI
- EX4550-32T-DC-AFO

EX4550 switches are available in two base models: a 32-port fiber-based model (EX4550-32F) that provides 32 fixed 1-gigabit or 10-gigabit SFP or SFP+ ports, and a 32-port copper-based model (EX4550-32T) that provides 32 fixed 100 megabit, 1GBASE-T, or 10GBASE-T Ethernet ports.

All EX4550 switches provide either AFI (airflow in) or AFO (airflow out) airflow, and they provide two module slots, each of which can house any one of the following optional modules:

- 8-port SFP+ expansion module
- 8-port 10GBASE-T expansion module
- 128-Gb Virtual Chassis module

EX4550 switches support 650-W hot-insertable and hot-removable field-replaceable unit (FRU) AC and DC power supplies.

EX4550 switches support the following optical transceivers:

- EX-SFP-1GE-LX
- EX-SFP-1GE-SX
- EX-SFP-1GE-T
- EX-SFP-10GE-DAC-1M
- EX-SFP-10GE-DAC-3M

- EX-SFP-10GE-DAC-5M
- EX-SFP-10GE-DAC-7M
- EX-SFP-10GE-ER
- EX-SFP-10GE-LR
- EX-SFP-10GE-LRM
- EX-SFP-10GE-SR
- EX-SFP-10GE-USR
- EX-SFP-10GE-ZR

[See [EX4550 Hardware Documentation](#).]

- **EX4550 Virtual Chassis module**—EX4550 switches support the 128-Gb Virtual Chassis module, which can connect EX4550 switches to EX4200 switches, EX4500 switches, or other EX4550 switches to form one unit that you can manage as a single Virtual Chassis. The Virtual Chassis module provides two dedicated Virtual Chassis ports (VCPs) that can be used to connect the switch to other Virtual Chassis member switches. However, it is not mandatory to install a Virtual Chassis module to connect an EX4550 switch in a Virtual Chassis configuration. [See [Virtual Chassis Module in EX4550 Switches](#).]
- **EX4550 8X1/10-Gb SFP+ expansion module**—You can install up to two 8X1/10-Gb SFP+ optional expansion modules in an EX4550 switch. Each expansion module provides eight SFP+ ports for connecting to core devices in a data center. You can install SFP or SFP+ transceivers in these ports. [See [Expansion Modules in EX4550 Switches](#).]
- **EX4550 8X10GBASE-T expansion module**—You can install up to two 8-port 10G BASE-T optional expansion modules in an EX4550 switch. Each expansion module provides eight ports for 100M/1G/10GBASE-T Ethernet connectors. [See [Expansion Modules in EX4550 Switches](#).]

Use only IEEE 802.3an specified cables to connect to on-board EX4550 copper 10GBASE-T Ethernet network ports and EX4550 copper 10GBASE-T expansion module ports. For more information about the 10GBASE-T cable specification, see [Network Cable Specifications for EX4550 Switches](#).

- **DOM support on Virtual Chassis ports**—EX3300, EX4200, and EX4500 switches now support DOM on Gigabit Ethernet and 10-Gigabit Ethernet VCPs. [See [Pluggable Transceivers Supported on EX3300 Switches](#), [Pluggable Transceivers Supported on EX4200 Switches](#), [Pluggable Transceivers Supported on EX4500 Switches](#), and the `show interfaces diagnostics optics` command.]

### [Access Control and Port Security](#)

- **Support for the Infranet Controller (IC) as an external captive-portal server**—If you have connected an EX Series switch to the Junos Pulse Access Control Service and you want to use the captive portal feature, use the Access Control Service as an external captive portal server. For accessing a protected network resource that is connected to the switch, a user must first sign in to the Access Control Service for authentication

and endpoint security checking. The captive portal redirects the user to a login page located on the Access Control Service. When the user logs in, the Access Control Service examines the endpoint for compliance with security policies. If the endpoint passes the security check, the user is authenticated and is granted access to the protected resource. [See [Understanding Centralized Network Access Control and EX Series Switches](#).]

- **Junos Pulse Access Control Service**—Junos Pulse Access Control Service eliminates the need for you to configure firewall filters on each EX Series switch. Instead, you define resource access policies centrally on the network access control (NAC) device. Resource access policies define which network resources are allowed or are denied for a user, based upon the user's role. The NAC device distributes these policies to all connected switches. The NAC device thus functions as a centralized policy management server. The switch converts resource access policies into filter definitions and applies these to the appropriate ports. [See [Junos Pulse Access Control Service and Using the EX Series Switch as an Infranet Enforcer](#).]

---

### Ethernet Switching and Spanning Trees

- **Filtering BPDUs without blocking the port**—The original BPDU port block disabled the port, preventing all traffic from forwarding through the port. The **drop** statement lets you filter BPDUs instead of disabling the port. The port drops only incompatible BPDUs that try to enter at the port. Any other ingress traffic continues to be forwarded when the port is active. [See [Understanding BPDU Protection for STP, RSTP, and MSTP on EX Series Switches](#) and the **drop** configuration statement.]
- **Increased number of RTGs on EX8200 Virtual Chassis**—RTGs provide redundancy in cases of link or line card failures. You can now configure up to 254 RTGs on either an EX8200 switch or an EX8200 Virtual Chassis. [See [Understanding Redundant Trunk Links on EX Series Switches](#).]
- **VSTP compatibility with Cisco PVST+**—When you configure VSTP using the **set protocol vstp vlan all** configuration mode command, VLAN ID 1 is now excluded, thus making Junos OS VSTP compatible with Cisco PVST+. To include VLAN ID 1 in the VSTP VLAN, you must now add it explicitly using the **set protocol vstp vlan 1** configuration mode command. [See [Understanding VSTP for EX Series Switches](#) and the **vlan (VSTP)** configuration statement.]

## Firewall Filters

---

- **EX8200 management counters for displaying policer billing information**—You can obtain policer statistics in EX8200 switches by using three global management counters. You can assign any number of ingress policers to each global management counter and obtain the policer statistics. The policer statistics for each global management counter are the aggregate of the policer statistics for all policers associated with that global management counter. [See [Understanding the Use of Policers in Firewall Filters](#), [Configuring Policers to Control Traffic Rates \(CLI Procedure\)](#), the `counter` configuration statement, and the `show firewall` command.]

## High Availability

---

- **NSR for IPv6 RIPng with BFD, IPv6 OSPFv3 with BFD, and IPv6 IS-IS with BFD support on EX3300 Virtual Chassis, EX4200 Virtual Chassis, and EX4500 Virtual Chassis**—NSR for IPv6 IS-IS with BFD, IPv6 OSPFv3 with BFD, and IPv6 RIPng with BFD is now supported on EX3300 Virtual Chassis, EX4200 Virtual Chassis, and EX4500 Virtual Chassis. You can now configure NSR to enable a transparent switchover between the master and backup Routing Engines without having to restart any of these protocols. [See [Understanding Nonstop Active Routing on EX Series Switches](#).]
- **NSR for OSPFv3 and RIPng with BFD support on EX4200 Virtual Chassis and EX4500 Virtual Chassis**—NSR for OSPFv3 and RIPng with BFD are now supported on EX4200 Virtual Chassis and EX4500 Virtual Chassis. You can now configure NSR to enable a transparent switchover between the master and backup Routing Engines without having to restart OSPFv3 and RIPng with BFD. [See [Understanding Nonstop Active Routing on EX Series Switches](#).]
- **NSR for PIM on EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, and EX6200 switches**—NSR for PIM is now supported on EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, and EX6200 switches. You can now configure NSR to enable a transparent switchover between the master and backup Routing Engines without having to restart PIM. [See [Understanding Nonstop Active Routing on EX Series Switches](#).]
- **NSSU support on EX3300 Virtual Chassis and EX6200 switches**—NSSU, which permits you to upgrade the software running on a switch or Virtual Chassis with minimal disruption to traffic, is now supported on EX3300 Virtual Chassis and on EX6200 switches. [See [Understanding Nonstop Software Upgrade on EX Series Switches](#).]

## Infrastructure

---

- **Automatic switch provisioning without manual intervention**—Automatic switch provisioning without manual intervention is now supported. When you physically connect a switch to the network and boot it with a default configuration, it attempts to upgrade software automatically and autoinstall a configuration file from the network. The switch uses information provided by a DHCP server to determine whether to perform these actions and to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration. [See [Understanding EZ Touchless Provisioning on EX Series Switches](#).]
- **Complete EX Series—specific configuration statement hierarchies**—Documentation is now provided that lists all supported and unsupported configuration statements in each configuration hierarchy level on EX Series switches. *Supported* statements are those that you can use to configure some aspect of a software feature on the switch. *Unsupported* statements are those that appear in the CLI on the switch, but that have no effect on switch operation if you configure them. [See the Configuration tab on [User Interfaces on EX Series Switches](#).]
- **Enhancements to the display of packet drop and error counters**—The software now provides additional debugging ability in retrieving packet drop information and errors. The following operational mode commands have been added or updated to enable you to retrieve packet drop information and errors: **show pfe statistics bridge**, **show pfe statistics errors**, and **show pfe statistics traffic**. [See the [show pfe statistics bridge](#), [show pfe statistics error](#), and [show pfe statistics traffic](#) commands.]
- **New software features for EX3300 switches**—Several new features that were introduced in earlier EX Series switches are now supported on EX3300 switches and EX3300 Virtual Chassis. [See [EX Series Switch Software Features Overview](#) and [EX Series Virtual Chassis Software Features Overview](#) for a list of supported features.]
- **Routing Engine SDK package**—The Routing Engine SDK (RE SDK) package is now supported on EX4200 standalone switches, EX4200 Virtual Chassis, EX4500 standalone switches, EX4500 Virtual Chassis, EX8200 standalone switches, EX8200 Virtual Chassis, and mixed EX4200 and EX4500 Virtual Chassis. The Junos Software Development Kit (SDK) enables partners of the Junos SDK program to build custom applications that run on Junos OS. The Routing Engine SDK enables developers to create applications to run on the control plane or Routing Engine. [See [Junos SDK](#).]

## Interfaces

---

- **Energy Efficient Ethernet**—Energy Efficient Ethernet (EEE) reduces the power consumption of BASE-T copper physical layers (PHYs) during periods of low link utilization. EEE, an Institute of Electrical and Electronics Engineers (IEEE) 802.3az standard, specifies a signaling protocol, Low Power Idle (LPI), to achieve the power-saving goal during the idle time of links. [See [Understanding How Energy Efficient Ethernet Reduces Power Consumption on Interfaces](#).]
- **Load balancing of multicast traffic over aggregated Ethernet interfaces on EX8200 switches**—You can now virtually aggregate four 10-gigabit links on EX8200 switches

to form a 40-gigabit point-to-point link channel for data. You can use the **show chassis multicast load-balance** command to see whether multicast load balancing is enabled and, if it is enabled, what the hash mode has been set to. [See [Understanding Multicast Load Balancing Over 10-Gigabit Links for Routed Multicast Traffic on Switches](#).]

### J-Web Interface

---

- **J-Web interface configuration for EX4550 32-F switch**—You can configure the EX4550 32-F switch in the J-Web interface. [See [User Interfaces on EX Series Switches](#).]

### Management and RMON

---

- **Juniper Networks enterprise-specific interface MIB enhancements**—The jnxIfTable in the Juniper Networks enterprise-specific interface MIB has been enhanced to display the count of the number of cyclic redundancy check (CRC) errors and frame check sequence (FCS) errors. [See [Juniper Networks Enterprise-Specific MIBs](#) and [Junos OS Enterprise MIBs](#).]

### MPLS

---

- **MPLS support on EX4500 standalone switches and EX4500 Virtual Chassis**—EX4500 standalone switches and EX4500 Virtual Chassis now support all MPLS features that are supported on EX8200 switches with the following exceptions:
  - MPLS is not supported in a mixed EX4200 and EX4500 Virtual Chassis. EX4500 switches support IP over MPLS only when the switches are configured to perform penultimate-hop popping (PHP). MPLS over RVIs, LSP statistics, unicast RPF statistics, MPLS CoS, traffic policing, DiffServ-aware LSPs, GRES, and equal-cost multipath (ECMP) are not supported.
  - EX4500 standalone switches and EX4500 Virtual Chassis support a maximum of 125 instances of Layer 2 VPN, Layer 3 VPN, or CCC connections; or a combination of these.
  - LSP ping and traceroute operations for CCCs, Layer 2 circuits, and Layer 2 VPNs are not supported.
  - An MPLS configuration that consists of a mix of EX8200 and EX4500 switches does not support VLAN CCCs.
  - VLAN CCCs require that the VLAN ID be the same at both ends of the connection. The VLAN ID translation feature is not supported.
  - The TTL of MPLS packets is not decremented in the ingress MPLS switch.
  - The pipe model of TTL handling is not supported on a Layer 3 VPN if an EX4500 switch is configured as the ingress PE switch.

[See [MPLS for EX Series Switches](#).]

- **MPLS support on EX8200 Virtual Chassis**—EX8200 Virtual Chassis now support MPLS on all member switches. [See [MPLS for EX Series Switches](#).]

- **MPLS protocol enhancements on EX8200 standalone switches and EX8200 Virtual Chassis**—EX8200 standalone switches and EX8200 Virtual Chassis now support the following protocols:

- BFD is a simple hello mechanism that detects failures in a network. Hello packets are sent at a specified, regular interval. A neighbor failure is assumed when the routing device stops receiving a reply from the neighbor after a specified interval. You can also use a ping operation to detect network failures. BFD is supported for LSPs (both RSVP and LDP), Layer 3 VPNs, and Multi-Gateway Multipath (MGMP) networks.
- The ping operation is now supported for LSPs and Layer 3 VPNs. Note that the processing resources required for BFD are much less than those required for a ping operation. In addition, BFD is capable of detecting data plane failure faster than the ping operation.
- For an MPLS-based Layer 3 VPN, you can perform a traceroute operation to display the route that packets take to a specified network host.

[See [MPLS for EX Series Switches](#).]

---

### Power over Ethernet (PoE)

- **LLDP PoE power negotiation**—EX2200, EX3300, EX4200 PX, EX4500, EX6200, and EX8200 switches now support LLDP PoE power negotiation. The switch can dynamically allocate PoE power to powered devices based on their needs and obtain the PoE priority value from powered devices using LLDP. [See [Understanding PoE on EX Series Switches](#).]

---

### Software Installation and Upgrade

- **Zero Touch Provisioning on EX3300 switches**—Zero Touch Provisioning is supported on EX3300 switches with Junos OS Release 12.2R5. Zero Touch Provisioning allows you to provision new Juniper Networks switches in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, it attempts to upgrade Junos OS automatically and autoinstall a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version. (“Zero Touch Provisioning” was previously named “EZ Touchless Provisioning.”) [See [Understanding EZ Touchless Provisioning](#). See also “Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches” on page 92.]

---

### Virtual Chassis

- **Dedicated Virtual Chassis port link aggregation on EX4550 switches**—The dedicated Virtual Chassis ports (VCPs) on EX4550 switches automatically form a link aggregation group (LAG) bundle when two or more dedicated VCPs are used to interconnect the same Virtual Chassis member switches. This feature became available with Junos OS



Release 12.2R4. The LAG provides more bandwidth than a single dedicated VCP can provide, and it provides VCP redundancy by load-balancing traffic across all available dedicated VCPs in the LAG. If one of the dedicated VCPs fails, the VCP traffic is automatically load-balanced across the remaining dedicated VCPs in the LAG. See also [“Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches”](#) on page 92.

- **Hardware rate limiting on XRE200 External Routing Engines**—Internal hardware rate limiting has been modified to help secure the XRE200 External Routing Engine control plane. [See [XRE200 External Routing Engine Documentation](#).]

#### Related Documentation

- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches](#) on page 49
- [Limitations in Junos OS Release 12.2 for EX Series Switches](#) on page 50
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches](#) on page 57
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches](#) on page 64
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches](#) on page 92
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches](#) on page 94

## Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches

This section lists the changes in default behavior and syntax in Junos OS Release 12.2 for EX Series switches.

### Ethernet Switching and Spanning Trees

---

- When you configure VSTP using the **set protocol vstp vlan all** configuration mode command, VLAN ID 1 is now excluded, thus making the Junos OS VSTP compatible with Cisco PVST+. To include VLAN ID 1 in the VSTP VLAN, you must now add it explicitly using the **set protocol vstp vlan 1** configuration mode command.

### High Availability

---

- Change in the automatically generated virtual-link-local-address for VRRP over IPv6—over IPv6 will be 0x02. This change makes the VRRP over IPv6 feature in Junos OS Releases 12.2R5, 12.3R3, and later releases inoperable with Junos OS Releases 12.2R1, 12.2R2, 12.2R3, 12.2R4, 12.3R1, and 12.3R2 if automatically generated virtual-link-local-address IDs are used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

### Infrastructure

---

- The switch behavior when you reboot EX Series switches that are in the default configuration has changed. When you physically connect a switch that is in the default configuration to the network and boot it, the switch attempts to upgrade software

automatically and autoinstall a configuration file from the network. The switch uses information provided by a DHCP server to determine whether to perform these actions and to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration.

The switch uses different DHCP options than those for previous releases to locate configuration files on the network when it is booted with a default configuration. For more information, see the documentation for automatic switch provisioning without manual intervention.

- The switch monitors available disk space in the **/var** partition every 10 minutes. If disk space in the **/var** partition is more than 75 percent of the partition space, the switch displays a yellow alarm. If disk space in the **/var** partition is more than 90 percent of the partition space, the switch displays both a yellow alarm and a red alarm. To avoid getting these warnings, use the **request system storage cleanup** command to clear disk space.
- The **show system services dhcp client** command on EX Series switches now includes decrementing-lease-time information to aid in debugging. The field **Lease expires in:** has been added to the command output.

#### Related Documentation

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Limitations in Junos OS Release 12.2 for EX Series Switches on page 50](#)
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches on page 57](#)
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches on page 64](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches on page 92](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches on page 94](#)

## Limitations in Junos OS Release 12.2 for EX Series Switches

This section lists the limitations in Junos OS Release 12.2 for EX Series switches. If the limitation is associated with an item in our bug database, the description is followed by the bug tracking number.

For the most complete and latest information about known Junos OS defects, use the [Juniper online Junos Problem Report Search application](#).

### Access Control and Port Security

---

- You cannot configure a security certificate to provide additional security for the connection between an EX Series switch and Junos Pulse Access Control Service; that is, the **ca-profile** and **server-certificate-subject** configuration statements are not supported. [This is a known software limitation.]
- On EX Series switches, you cannot configure 802.1X authentication on redundant trunk groups (RTGs). [This is a known software limitation.]

## Ethernet Switching and Spanning Trees

- If the bridge priority of a VSTP root bridge is changed such that this bridge becomes a nonroot bridge, the transition might take more than 2 minutes, and you might see a loop during the transition. [PR/661691: This is a known software limitation.]
- On EX Series switches, only dynamically learned routes can be imported from one routing table group to another. [This is a known software limitation.]

## Firewall Filters

- On EX3200 and EX4200 switches, when a very large number of firewall filters are included in the configuration, it might take a long time, possibly a few minutes, for the egress filter rules to be installed. [PR/468806: This is a known software limitation.]
- On EX3300 switches, if you add and delete filters with a large number of terms (on the order of 1000 or more) in the same commit operation, not all the filters are installed. As a workaround, add filters in one commit operation, and delete filters in a separate commit operation. [PR/581982: This is a known software limitation.]
- On EX8200 switches, if you configure an implicit or explicit discard action as the last term in an IPv6 firewall filter on a loopback (lo0) interface, all the control traffic from the loopback interface is dropped. To prevent this, you must configure an explicit **accept** action. [This is a known software limitation.]

## Hardware

- On 40-port SFP+ line cards for EX8200 switches, the LEDs on the left of the network ports do not blink to indicate that there is link activity if you set the speed of the network ports to 10/100/1000 Mbps. However, if you set the speed to 10 Gbps, the LEDs blink. [PR/502178: This is a known limitation.]
- The [Uplink Modules in EX3200 Switches](#) topic notes the following behavior for the SFP uplink module, which provides four ports for 1-gigabit SFP transceivers: “On an EX3200 switch, if you install a transceiver in an SFP uplink module, a corresponding network port from the last four built-in ports is disabled. For example, if you install an SFP transceiver in port 2 on the uplink module (ge-0/1/2) on 24-port models, then ge-0/0/22 is disabled. The disabled port is not listed in the output of **show interface** commands.”

Another note on the same page describes similar behavior for the SFP+ uplink module: “On an EX3200 switch, if you install a transceiver in an SFP+ uplink module when the uplink module is operating in 1-gigabit mode, a corresponding network port from the last four built-in ports is disabled. For example, if you install an SFP transceiver in port 2 on the uplink module (ge-0/1/2), then ge-0/0/22 is disabled. The disabled port is not listed in the output of **show interfaces** commands.”

However, in both cases what actually occurs is that when you install the SFP uplink module or explicitly configure the mode on an SFP+ uplink module to 1-gigabit operating mode and do not reboot the switch, the last four built-in ports on the switch are disabled. If transceivers are installed in the uplink module, the corresponding built-in network ports are not displayed in the output of **show interfaces** commands. The

workaround is to move all four links to the uplink module, or to reboot the switch for correct initialization of the ports.

[PR/686467: This is a known limitation.]

- You cannot connect EX2200-12P switches to some vendors' prestandard IP phones with a straight cable. As a workaround, use a crossover cable. [PR/726929: This is a known limitation.]

### High Availability

---

- You cannot verify that NSB is synchronizing Layer 2 protocol information with the backup Routing Engine even when NSB is properly configured. [PR/701495: This is a known software limitation.]
- On EX Series Virtual Chassis running Junos OS Release 11.2 or earlier, the same MAC address might be assigned to multiple Layer 2 interfaces and aggregated Ethernet interfaces on different member switches. This is an expected behavior: you cannot assign a unique MAC address to each interface when the Virtual Chassis is running Junos OS Release 11.2 or earlier. However, starting in Junos OS Release 11.3, you can assign unique MAC addresses to these Virtual Chassis interfaces.

If you use NSSU to upgrade a Virtual Chassis from Junos OS Release 11.2 or earlier to Junos OS Release 11.3 or later, you might see the same MAC address assigned to multiple interfaces on different member switches. To ensure that the interfaces have unique MAC addresses, either perform the upgrade without using NSSU or reboot the Virtual Chassis after you perform the upgrade with NSSU.

[PR/775203: This is a known software limitation.]

### Infrastructure

---

- On EX Series switches, the **show snmp mib walk etherMIB** command does not display any output, even though the etherMIB is supported. This occurs because the values are not populated at the module level—they are populated at the table level only. You can issue the **show snmp mib walk dot3StatsTable**, **show snmp mib walk dot3PauseTable**, and **show snmp mib walk dot3ControlTable** commands to display the output at the table level. [This is a known software limitation.]
- Momentary loss of an inter-Routing Engine IPC message might trigger an alarm that displays the message **Loss of communication with Backup RE**. However, no functionality is affected. [PR/477943: This is a known software limitation.]
- Routing between virtual routing instances for local direct routes is not supported. [PR/490932: This is a known software limitation.]
- On EX4500 switches, the maintenance menu is not disabled even if you include the **lcd maintenance-menu disable** statement in the configuration. [PR/551546: This is a known software limitation.]
- When you enable the **filter-id** attribute on the RADIUS server for a particular client, none of the required 802.1X authentication rules are installed in the IPv6 database. Therefore, IPv6 traffic on the authenticated interface is not filtered; only IPv4 traffic is filtered on that interface. [PR/560381: This is a known software limitation.]

- On EX8200 switches, if OAM link-fault management (LFM) is configured on a member of a VLAN on which Q-in-Q tunneling is also enabled, OAM PDUs cannot be transmitted to the Routing Engine. [PR/583053: This is a known software limitation.]
- When you reconfigure the MTU value of a next hop more than eight times without restarting the switch, the interface uses the maximum value of the eight previously configured values as the next MTU value. [PR/590106: This is a known software limitation.]
- On EX8208 and EX8216 switches that have two Routing Engines, one Routing Engine cannot be running Junos OS Release 10.4 or later while the other one is running Junos OS Release 10.3 or earlier. Ensure that both Routing Engines in a single switch run either Junos OS Release 10.4 or later or Junos OS Release 10.3 or earlier. [PR/604378: This is a known software limitation.]
- When you configure a static route that has two multihop paths, BFD might become unstable, and the routing protocol process (rpd) might crash. [PR/701966: This is a known software limitation.]
- On EX6210 and EX8200 switches that have two Routing Engines, and on EX8200 Virtual Chassis that have two XRE200 External Routing Engine modules, you cannot issue the **commit synchronize** command from the J-Web interface. As a workaround, issue this command from the CLI. [This is a known software limitation.]
- If the accounting server is not available, you might experience trouble viewing system information on EX4200 switches. If you attempt to execute CLI commands related to system options, the following error message might be displayed: **error communicating with fpc0**. This error is a result of the common command forwarding used by EX4200 switches to gather information about other members in a Virtual Chassis. When you issue commands related to system options through the CLI, a new management process (mgd) is initiated. The management process on one Virtual Chassis member opens a connection to a management process on another member, logs in, and extracts the information. Because the EX4200 switch is a Virtual Chassis, this sequence of events occurs even on a standalone EX4200 switch. The connection from one mgd process to another is treated as a login event. If system accounting is configured for login events, the switch attempts to connect to the accounting server before executing the CLI command. If the accounting server is not available, the connection times out. As a workaround, either ensure that the accounting server is reachable or disable the configuration of system accounting for login events. [This is a known software limitation.]

## Interfaces

- EX Series switches do not support IPv6 interface statistics. Therefore, all values in the output of the **show snmp mib walk ipv6IfStatsTable** command always display a count of 0. [PR/480651: This is a known software limitation.]
- On EX8216 switches, a link might go down momentarily when an interface is added to a LAG. [PR/510176: This is a known software limitation.]
- On EX Series switches, if you clear LAG interface statistics while the LAG is down, then bring up the LAG and pass traffic without checking for statistics, and finally bring the LAG interface down and check interface statistics again, the statistics might be

inaccurate. As a workaround, use the **show interfaces interface-name** command to check LAG interface statistics before bringing down the interface. [PR/542018: This is a known software limitation.]

- Power over Ethernet (PoE) and Power over Ethernet Plus (PoE+) cannot be configured for EX8200 member switches in an EX8200 Virtual Chassis using the XRE200 External Routing Engine.

If you have not cabled the Virtual Chassis, configure PoE or PoE+ on each EX8200 member switch before cabling the Virtual Chassis. See *Configuring PoE (CLI Procedure)*.

To configure PoE and PoE+ on an EX8200 member switch in an operational EX8200 Virtual Chassis:

1. Power off the EX8200 member switch. See *Powering Off an EX8200 Switch*.
2. Uncable the switch from the Virtual Chassis.
3. Power on the switch. See *Powering On an EX8200 Switch*.
4. Log in to the switch. See *Connecting an EX Series Switch to a Management Console*.
5. Configure PoE. See *Configuring PoE (CLI Procedure)*.
6. Cable the EX8200 member switch back into the EX8200 Virtual Chassis. See *Connecting an EX8200 Switch to an XRE200 External Routing Engine*.

[This is a known software limitation.]

### J-Web Interface

---

- In the J-Web interface, you cannot commit some configuration changes in the Ports Configuration page or 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 interface.

[PR/400814: This is a known software limitation.]

- In the J-Web interface, the Ethernet Switching Monitor page (Monitor > Switching > Ethernet Switching) might not display monitoring details if the switch has more than 13,000 MAC entries. [PR/425693: This is a known software limitation.]
- In the J-Web interface for EX4500 switches, the Port Configuration page (Configure > Interfaces > Ports), the Port Security Configuration page (Configure > Security > Port Security), and the Filters Configuration page (Configure > Security > Filters) display features that are not supported on EX4500 switches. [PR/525671]
- When you use an HTTPS connection in the Microsoft Internet Explorer browser to save a report from the following pages in the J-Web interface, the error message **Internet Explorer was not able to open the Internet site** is displayed on the following pages:

- Files page (Maintain > Files)
- History page (Maintain > Config Management > History)
- Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
- Static Routing page (Monitor > Routing > Route Information)
- Support Information page (Maintain > Customer Support > Support Information)
- View Events page (Monitor > Events and Alarms > View Events)

[PR/542887]

- If you insert four or more EX8200-40XS line cards in an EX8208 or EX8216 switch, the Support Information page (Maintain > Customer Support > Support Information) in the J-Web interface might fail to load because the configuration might be larger than the maximum size of 5 MB. The error message **Configuration too large to handle** is displayed. [PR/552549: This is a known software limitation.]
- In the J-Web interface, you cannot configure interface ranges and interface groups. [This issue was being tracked by PR/600559.]
- The J-Web interface does not support role-based access control; it supports users in the super-user authorization class only. So a user who is not in the super-user class, such as a user with view-only permission, is able to launch the J-Web interface and is allowed to configure everything, but the configuration fails on the switch, and the switch displays access permission errors. [PR/604595: This is a known software limitation.]

### Layer 2 and Layer 3 Protocols

---

- On EX 3200 and EX4200 switches, MPLS on Layer 3 tagged subinterfaces and routed VLAN interfaces (RVIs) is not supported, even though you can commit a configuration that enables these features. [PR/612434: This is a known software limitation.]

### Management and RMON

---

- 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. [PR/441789: This is a known software limitation.]
- When MVRP is configured on a trunk interface, you cannot configure connectivity fault management (CFM) on that interface. [PR/540218: This is a known software limitation.]
- EX Series switches do not support sFlow monitoring technology on equal-cost multipath (ECMP) flow-based forwarding. Therefore, in an ECMP scenario, information about the egress interface that the Packet Forwarding Engine uses cannot be fetched by sFlow monitoring technology. [This is a known software limitation.]

## Software Installation and Upgrade

---

- Do not manually configure the **autoinstallation** statement at the **[edit system]** configuration hierarchy level. If you configure this statement, the switch loses network connectivity and cannot respond either to ping requests or to ARP requests. The initial default configuration that is present on the switch does include the **autoinstallation** statement at the **[edit system]** hierarchy level. The reason is that if the switch cannot find any configuration file stored in flash memory, it must try to obtain one from a TFTP server. However, the first time you commit a configuration on the switch, a configuration file is created in flash memory, and the statement is deleted from the configuration because it is no longer needed. If you do manually configure the **autoinstallation** statement and your switch loses network connectivity, delete the **autoinstallation** statement from the configuration and commit the configuration. [PR/610816: This is a known software limitation.]

## Virtual Chassis

---

- A standalone EX4500 switch with its PIC mode set to **virtual-chassis** has less bandwidth available for network ports than an EX4500 switch with its PIC mode set to **intraconnect**. The network ports on a standalone EX4500 switch with a **virtual-chassis** PIC mode setting often do not achieve line-rate performance.

The PIC mode on an EX4500 switch can be set to **virtual-chassis** if:

- The switch was ordered with a Virtual Chassis module installed and thus has its PIC mode set to **virtual-chassis** by default.
- You entered the **request chassis pic-mode virtual-chassis** operational mode command to configure the switch as a member of a Virtual Chassis.

You can check the PIC mode for your EX4500 switch that has a Virtual Chassis module installed by entering the **show chassis pic-mode** command.

You must always set the PIC mode on a standalone EX4500 switch to **intraconnect**. Set the PIC mode to **intraconnect** by entering the **request chassis pic-mode intraconnect** operational mode command.

[This is a known software limitation.]

- The automatic software update feature is not supported on EX4500 switches that are members of a Virtual Chassis. [PR/541084: This is a known software limitation.]
- When an EX4500 switch becomes a member of a Virtual Chassis, it is assigned a member ID. If that member ID is a nonzero value, then if the software on that member switch is downgraded to a software image that does not support Virtual Chassis, you cannot change the member ID to 0. A standalone EX4500 switch must have a member ID of 0. The workaround is to convert the EX4500 Virtual Chassis member switch to a standalone EX4500 switch before downgrading the software to an earlier release, as follows:
  1. Disconnect all Virtual Chassis cables from the member to be downgraded.
  2. Convert the member switch to a standalone EX4500 switch by issuing the **request virtual-chassis reactivate** command.



3. Renumber the member ID of the standalone switch to 0 by issuing the **request virtual-chassis renumber** command.
4. Downgrade the software to the earlier release.

[PR/547590: This is a known software limitation.]

- When you add a new member switch to an existing EX4200 Virtual Chassis, EX4500 Virtual Chassis, or mixed EX4200 and EX4500 Virtual Chassis in a ring topology, a member switch that was already part of the Virtual Chassis might become nonoperational for several seconds. The member switch returns to the operational state with no user intervention. Network traffic to the member switch is dropped during the downtime. To avoid this issue, follow this procedure:
  1. Cable one dedicated or user-configured Virtual Chassis port (VCP) on the new member switch to the existing Virtual Chassis.
  2. Power on the new member switch.
  3. Wait for the new switch to become operational in the Virtual Chassis. Monitor the **show virtual-chassis** command output to confirm the new switch is recognized by the Virtual Chassis and is in the Prsnt state.
  4. Cable the other dedicated or user-configured VCP on the new member switch to the Virtual Chassis.

[PR/591404: This is a known software limitation.]

- On EX4550 Virtual Chassis, if you configure more than 2036 routed VLAN interfaces (RVIs), ping operations might fail. [PR/791821: This is a known software limitation.]
- On EX4200 Virtual Chassis, when you perform an NSSU upgrade from Junos OS Release 12.1R5 to Junos OS Release 12.2R3, a vmcore file might be created in linecard members. [PR/844519]

#### Related Documentation

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches on page 49](#)
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches on page 57](#)
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches on page 64](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches on page 92](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches on page 94](#)

## Outstanding Issues in Junos OS Release 12.2 for EX Series Switches

The following are outstanding issues in Junos OS Release 12.2R9 for EX Series switches. The identifier following the description is the tracking number in our bug database.

For the most complete and latest information about known Junos OS defects, use the [Juniper online Junos Problem Report Search application](#).



**NOTE:** Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in [“Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers” on page 171](#).

---

### Access Control and Port Security

---

- On aggregated Ethernet (ae) interfaces, LLDP might not work. [PR/781814]
- On EX Series switches with DHCP snooping enabled, the DHCP reply packets without any DHCP options (BOOTP reply packets) might be dropped. [PR/925506]

---

### Hardware

---

- On EX4200 switches, the serial number displayed in the log output does not include the leading 0 that appears on the switch label. This problem is just a display issue—there is no functional impact. [PR/815950]
- On EX Series switches, an SFP might stop working unexpectedly with i2c errors and the switch might not recognize the SFP in its existing port. [PR/939041]

---

### Infrastructure

---

- When multicast traffic is transiting an EX8200 switch, kernel panic might occur on a new master Routing Engine, and the string **rn\_clone\_unwire parent unreferenced** might appear during NSSU or after multiple GRES operations. [PR/734295]
- When a hostname is added as an NTP server, it resolves to an IP address before it is added to the configuration. When you use a public NTP server, the hostname might resolve to different IP addresses. If the resolved IP address becomes unreachable for any reason, the switch cannot reach the NTP server. To leverage public pool entities, this has been modified so that a hostname is accepted as a string without DNS resolution. [PR/755591]
- On an EX4550 Virtual Chassis, the **show chassis environment power-supply-unit** CLI command does not show the power supply status of all the member interfaces. Use the **show chassis hardware** CLI command to see the status of the power supplies in various member interfaces. [PR/817397]
- On EX Series switches, when a VRRP switchover is initiated, the backup switch does not send out an unsolicited Neighbor Advertisement message. [PR/824465]
- If you delete an IPv6 configuration on an RVI, ARP requests might not be trapped to the CPU and are not resolved. As a workaround, delete the RVI and then reconfigure it, or reboot the switch. [PR/826862]
- On EX4500 switches, a firewall filter with the **family** option set to **ethernet-switching** and configured for IPv4 might block specific IPv6 traffic. As a workaround, set the **ether-type** option to **ipv4** in the filter. [PR/843336]

- On EX8200 Virtual Chassis, a disabled RVI might send gratuitous ARP requests. [PR/848852]
- On EX Series switches, if the NTP server is not a stratum 1 server, the NTP synchronization process might fail. [PR/864223]
- On EX4550 switches, the log message **PFC is supported only on 10G interfaces** is generated over and over again in logs. [PR/880571]
- On EX8200 switches, if a line card is rebooted, traffic flow that is not related to the rebooted line card might get dropped. [PR/882257]
- On EX4550 switches, the link between the EX4550 and a third-party switch might not come up in 1-gigabit mode using an SFP fiber module. There is no issue with bringing up links when 10-gigabit SFP+ fiber modules or 1-gigabit SFP copper modules are used. [PR/886563]
- On EX4500 switches, the TLV type 314 is sent as a notification of the DCBX state of a port. In a link-flap scenario, the kernel might send a DCBX PFC state TLV to the Packet Forwarding Engine even if there is no change in the DCBX state. Also, the kernel might then sync this state to the backup Routing Engine. On the backup Routing Engine, this message is not processed, and the system shows an Unknown TLV type 314 error. The message itself is harmless, but it fills up the logs unnecessarily. [PR/893802]
- In EX4200 Virtual Chassis, a member of the Virtual Chassis might reboot and create a pfem core file. [PR/912889]
- On EX3200, EX3300, and EX4200 switches, if multicast traffic is bursty or cyclical with no traffic for continuous 30-second periods, the multicast keepalive timer might age out and then delete that particular route.

As a workaround, use one of the two following options:

- Set a large timeout value for multicast forwarding cache entries using the **set routing-options multicast forwarding-cache timeout** command if the traffic pattern is as described above.
- Using a script, issue the **show multicast route** command continuously every 25 seconds.

[PR/937695]

- On EX4500 or EX4550 switches, the software forwarding infrastructure daemon (sfid) might continuously create core files, causing interruptions in traffic, because packets are erroneously freed twice. A possible trigger is the handling of Layer 2 protocol tunneling packets. [PR/941482]
- On EX Series switches configured for accounting based on 802.1X RADIUS, if the RADIUS server is enabled with the *User-Name* attribute and a new user name is used to send account information, the switches might ignore this attribute and not send accounting information with the authentication user name. [PR/950562]
- On EX Series switches with Protocol Independent Multicast (PIM) configured, when the upstream interface on a rendezvous point (RP) changes between a Layer 3 interface and the PIM de-encapsulation interface for the multicast route, the earlier route entry might be deleted twice, causing a loss of multicast traffic on the RP. [PR/982883]

## Interfaces

---

- When you disable a static LAG on an aggregated Ethernet (ae) interface, Ethernet ring protection traffic traveling in one direction might be lost for 3 to 5 seconds, and traffic traveling in the other direction might contain extra packets. [PR/703091]

## J-Web Interface

---

- On EX Series switches and on SRX3400, SRX3600, SRX5600, and SRX5800 Series Services Gateways, when you use the Microsoft Internet Explorer browser to open reports from the following pages in the J-Web interface, the reports open in the same browser session:

- Files page (Maintain > Files)
- History page (Maintain > Config Management > History)
- Port Troubleshooting page (Troubleshoot > Troubleshoot > Troubleshoot Port)
- Static Routing page (Monitor > Routing > Route Information)
- Support Information page (Maintain > Customer Support > Support Information)
- View Events page (Monitor > Events and Alarms > View Events)

[PR/433883]

- In the J-Web interface, in the Port Security Configuration page, you are required to configure the **action** option when you configure the **MAC limit** option even though configuring an action value is not mandatory in the CLI. [PR/434836]
- In the J-Web interface on EX4200 switches; SRX100, SRX210, SRX240, and SRX650 Security Gateways; and all J Series devices, if you try to change the position of columns using the drag-and-drop method, only the column header moves to the new position instead of the entire column in the OSPF Global Settings table in the OSPF Configuration page, the Global Information table in the BGP Configuration page, or the Add Interface window in the LACP (Link Aggregation Control Protocol) Configuration page. [PR/465030]
- If you configure an IPv6 address for a VLAN in the J-Web interface, you cannot then edit the VLAN configuration. As a workaround, manually refresh the page and then edit the VLAN configuration. [PR/466633]
- When a large number of static routes are configured and you have navigated to pages other than page 1 in the Route Information table in the Static Routing monitoring page in the J-Web interface (Monitor > Routing > Route Information), changing the Route Table to query other routes refreshes the page but does not return to page 1. For example, if you run a query from page 3 and the new query returns very few results, the Results table continues to display page 3 and shows no results. To view the results, navigate to page 1 manually. [PR/476338]
- When you open a J-Web interface session using HTTPS, enter a username and password, and then click the Login button, the J-Web interface takes 20 seconds longer to launch and load the Dashboard page than it does if you use HTTP. [PR/549934]

- In the J-Web interface, you cannot upload a software package using the HTTPS protocol. As a workaround, use either the HTTP protocol or the CLI. [PR/562560]
- If you have accessed the J-Web interface using an HTTPS connection through the Microsoft Internet Explorer Web browser, you might not be able to download and save reports from some pages on the Monitor, Maintain, and Troubleshoot tabs. Some affected pages are at these locations:
  - Maintain > Files > Log Files > Download
  - Maintain > Config Management > History
  - Maintain > Customer Support > Support Information > Generate Report
  - Troubleshoot > Troubleshoot Port > Generate Report
  - Monitor > Events and Alarms > View Events > Generate Report
  - Monitor > Routing > Route Information > Generate Report

As a workaround, use the Mozilla Firefox Web browser to download and save reports using an HTTPS connection. [PR/566581]

- If you access the J-Web interface using the Microsoft Internet Web browser version 7, in the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR/603669]
- In the Process Details page (Monitor > System View > Process Details) of the J-Web interface, there are multiple entries for a few processes that do not impact any functionality. [PR/661704]
- In the J-Web interface, the Next Hop column in the Static Routing page (Monitor > Routing > Route Information) displays only the interface address, and the corresponding IP address is missing. The title of the first column displays Static Route Address instead of Destination Address. As a workaround, use the CLI to execute the **show route detail** command to fetch the corresponding next-hop interface IP address. [PR/684552]
- In the J-Web interface, HTTPS access might work with an invalid certificate. As a workaround, after you change the certificate, issue the **restart web-management** command to restart the J-Web interface. [PR/700135]
- In the J-Web interface, you cannot configure a large VLAN range. For example, you cannot configure the range to be 1 through 4093. [PR/700873]
- On EX4500 Virtual Chassis, if you use the CLI to switch from **virtual-chassis** mode to **intraconnect** mode, the J-Web interface dashboard might not list all the Virtual Chassis hardware components, and the images of the master and backup members of the Virtual Chassis might not be visible after an autorefresh occurs. The J-Web interface dashboard also might not list the vcp-0 and vcp-1 Virtual Chassis ports (VCPs) in the rear view of an EX4200 switch (in the linecard role) that is part of a mixed EX4200 and EX4500 Virtual Chassis. [PR/702924]

- On EX2200-C switches, if you change the media type and commit the change, the Ports Configuration page (**Configure > Interfaces > Ports**) might not list the uplink port. [PR/742847]
- In the J-Web interface, you cannot configure the TCP fragment flag for a firewall filter on the Filters Configuration page (**Configure > Security > Filters**). [PR/756241]
- If you create a Virtual Chassis by adding a switch to a standalone switch on which a J-Web session is already open, the chassis viewer might be aligned incorrectly on the dashboard. As a workaround, manually refresh the J-Web interface session. [PR/756711]
- In the J-Web interface, you cannot delete a term from a firewall filter and simultaneously add a new term to that filter in the Filters Configuration page (**Configure > Security > Filters**). [PR/769534]
- After you remove or reboot a Virtual Chassis member (either the backup or a member in the linecard role), when you click other members in the J-Web interface, the chassis view for those members might not expand, and the dashboard might log the following error: **stackImg is null or not an object**. [PR/771415]
- On EX Series Virtual Chassis that have more than five members, logging in to the J-Web interface dashboard might take more than 30 seconds. [PR/785300]
- On EX8200 Virtual Chassis, if you are using the Virtual Chassis wizard in the J-Web interface in the Mozilla Firefox version 3.x browser and have selected more than six port pairs from the same member to convert from VCPs to network ports, the wizard might display incorrect port conversion status. Also, if you double-click **Next** after deleting an active member in the Members page, the J-Web interface might stop working. [PR/796584]
- In the J-Web interface, you cannot configure OSPFv3 by using the point-and-click function (**Configure > Point&Click > Protocols > Configure > Ospf3**). As a workaround, configure OSPFv3 options by using the CLI. You can then view and edit the OSPFv3 parameters by using the point-and-click function in the J-Web interface. [PR/857540]

### Routing Protocols

---

- On an EX4500 switch, if a session with an unconfigured peer is initiated, and the peer AS is a member of a confederation, then an RPD core file might be created. As a workaround, use an explicitly configured peer for peers in the confederation ASes. [PR/963565]

### Software Upgrade and Installation

---

- On EX8200 Virtual Chassis, when you initiate an NSSU operation to upgrade from Junos OS Release 12.2R6 to Release 12.3R4, multiple pfem core files might be created on some or any of the EX8200 member switches. [PR/917863]

### Spanning-Tree Protocols

---

- On EX4550 Virtual Chassis with xSTP (RSTP, VSTP, or MSTP) enabled, multiple xSTP-enabled ports might go into an STP Disabled state on the Packet Forwarding Engine. The overlapping of STP identifiers might cause traffic to be dropped on these ports when this issue occurs. [PR/980551]

### Virtual Chassis

---

- On EX8200 Virtual Chassis, when you perform an snmpwalk operation on the jnxPsmMIB, the output shows details only for the power supplies on a single linecard member. [PR/689656]
- When you remove the hard drive from an XRE200 External Routing Engine, an SNMP trap and a system alarm are not generated. [PR/710213]
- On EX4200 switches, if you configure a physical interface on the master switch as a member of an interface range and associate that interface with a VLAN, then delete the interface from the interface range, the interface is not removed from the VLAN. [PR/811773]
- On EX2200 Virtual Chassis, when there are multiple equal-cost paths, the CLI command **show virtual chassis vc-path source-interface *interface-name* destination-interface *interface-name*** displays the first discovered shortest path, even though traffic might be flowing in an alternate path. As a workaround, discover the actual traffic path by checking traffic statistics on vc-ports using the CLI command **show virtual-chassis vc-port statistics**. [PR/829752]
- On EX4550 Virtual Chassis, broadcast traffic is not flooded to all Virtual Chassis members. [PR/850244]
- On EX8200 Virtual Chassis, when an NSSU is initiated to upgrade to Junos OS Release 12.3R5, multiple pfem core files might be created on some member switches. [PR/917863]

#### Related Documentation

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches on page 49](#)

- [Limitations in Junos OS Release 12.2 for EX Series Switches on page 50](#)
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches on page 64](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches on page 92](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches on page 94](#)

## Resolved Issues in Junos OS Release 12.2 for EX Series Switches

The following are the issues that have been resolved in Junos OS Release 12.2 for EX Series switches. The identifier following the descriptions is the tracking number in our bug database.

For the most complete and latest information about known Junos OS defects, use the [Juniper online Junos Problem Report Search application](#).



NOTE: Other software issues that are common to both EX Series switches and M, MX, and T Series routers are listed in [“Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers” on page 171](#).

- [Issues Resolved in Release 12.2R1 on page 64](#)
- [Issues Resolved in Release 12.2R2 on page 73](#)
- [Issues Resolved in Release 12.2R3 on page 77](#)
- [Issues Resolved in Release 12.2R4 on page 78](#)
- [Issues Resolved in Release 12.2R5 on page 80](#)
- [Issues Resolved in Release 12.2R6 on page 83](#)
- [Issues Resolved in Release 12.2R7 on page 85](#)
- [Issues Resolved in Release 12.2R8 on page 87](#)
- [Issues Resolved in Release 12.2R9 on page 89](#)

### Issues Resolved in Release 12.2R1

---

The following issues have been resolved since Junos OS Release 12.1. The identifier following the description is the tracking number in our bug database.

#### ***Access Control and Port Security***

- You cannot configure the level for storm control. [PR/734307: This issue has been resolved.]
- EX3200 switches might repeatedly create 802.1X core files. As a workaround, if access accounting is enabled, disable it by issuing the **deactivate access profile *profile-name* accounting** configuration mode command. [PR/739921: This issue has been resolved.]
- When you configure MVRP, the LLDP process might create a core file as the result of a memory leak. [PR/740793: This issue has been resolved.]



- If you enable 802.1X with MAC RADIUS authentication, that is, by including the **mac-radius** statement in the configuration, the authentication manager process (authd) might reach a memory limit when there are approximately 250 users. As a workaround, reset the authd process when it reaches 85 percent of its RLIMIT\_DATA value (that is, 85 percent of 130 MB). To check the amount of memory being used by the authd process, use the **show system processes extensive** operational mode command. [PR/783363: This issue has been resolved.]
- When access configuration is not required and the guest VLAN feature is configured, supplicants might not be authenticated using the guest VLAN, and they might remain in the connecting state. [PR/783606: This issue has been resolved.]
- DHCP snooping might prevent DHCP Inform ACK packets from passing to the client. [PR/787161: This issue has been resolved.]
- If you configure 802.1X (dot1X) with static MAC bypass and a new host is added to the exclusion list, the MAC addresses of existing hosts that have already been successfully authenticated by static MAC bypass might move to an incorrect VLAN. [PR/787679: This issue has been resolved.]

### ***Converged Networks (LAN and SAN)***

- On EX4500 switches, the DCBX protocol does not work. [PR/795835: This issue has been resolved.]

### ***Ethernet Switching and Spanning Trees***

- When you enable Q-in-Q tunneling and MLD snooping, no snooping database is present on the switch. [PR/693224: This issue has been resolved.]
- On EX Series switches, during the MAC learning period, excessive log messages similar to **MRVL-L2:mrvl\_fdb\_mac\_entry\_uc\_set()** might be displayed. [PR/695200: This issue has been resolved.]
- On XRE200 External Routing Engines, Layer 3 traffic on an RVI might fail. [PR/732237: This issue has been resolved.]
- When you configure an IPv6 address on a VLAN interface that is down, the IPv6 address might go into the Tentative state. [PR/733651: This issue has been resolved.]
- When using VSTP, if you try to enable all VLANs on a physical interface that is a member of all the VLANs, a configuration error might be displayed. For more information, see [“Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches” on page 94](#). [PR/736488: This issue has been resolved.]
- If a VLAN change occurs quickly, the client might not be able to obtain an IP address. [PR/746479: This issue has been resolved.]
- When VRRP is running between two EX8200 switches on a VLAN, after a master switchover, both switches might act as master. [PR/752868: This issue has been resolved.]

### **Firewall Filters**

- If multiple firewall rules are being programmed into the switch hardware simultaneously, a Packet Forwarding Engine (pfem) core file might be created. [PR/746337: This issue has been resolved.]

### **Hardware**

- After you have disabled the LCD Maintenance Menu and rebooted the switch, the EZSetup option might be available. [PR/707279: This issue has been resolved.]
- The EZSetup option is available on the LCD Maintenance Menu regardless of the factory default status of the switch. [PR/736411: This issue has been resolved.]
- On EX3300 switches, power supply failure errors might occur. To circumvent this problem, a software workaround has been provided. The software reads the power supply bit multiple times before it declares the power supply module to be down. [PR/743115: This issue has been resolved.]

### **High Availability**

- No CLI command is available to verify that NSB is enabled. To do this, you can now use the **show ethernet-switching task replication** command. [PR/613452: This issue has been resolved.]
- If you perform an NSSU operation that includes the **reboot** option, some traffic loss might occur. [PR/717662: This issue has been resolved.]
- On an XRE200 External Routing Engine, when you perform an NSSU operation that includes the **reboot** option, the physical link might flap, which causes traffic loss and protocol flapping. [PR/718472: This issue has been resolved.]
- When NSB is enabled on a switch, if you issue the **show spanning-tree interface msti msti-id** command on the backup Routing Engine, no output is displayed. [PR/732676: This issue has been resolved.]
- When you configure NSR for IP multicast, RPD on the backup switch might create a core file. [PR/734769: This issue has been resolved.]
- During a GRES operation, ICMP packets might be dropped. [PR/737168: This issue has been resolved.]
- After a GRES operation with nonstop bridging, the MSTP port boundary status might be displayed incorrectly. [PR/737179: This issue has been resolved.]

### **Infrastructure**

- If you enable gratuitous ARP by including the **gratuitous-arp-reply**, **no-gratuitous-arp-reply**, or **no-gratuitous-arp-request** statement in the configuration, the switch might process gratuitous ARP packets incorrectly. [PR/518948: This issue has been resolved.]
- Rate limiting for management traffic (namely, FTP, SSH, and Telnet) arriving on network ports causes file transfer speeds to be slow. [PR/691250: This issue has been resolved.]

- In some cases, broadcast traffic that is received on the management port (me0) is broadcast to other subnets on the switch. [PR/705584: This issue has been resolved.]
- In previous releases, typing the Alt-break sequence on the console put the console interface in debugger mode (the db> prompt). You can now configure the **system no-debugger-on-alt-break** statement to disable the Alt-break sequence on the serial console. [PR/717491: This issue has been resolved.]
- When the switch power-cycles ungracefully, the contents of flash memory and the switch's file system might become out of sync. [PR/719101: This issue has been resolved.]
- The **allow-configuration-regexps** statement at the **[edit system login class]** hierarchy level does not work exactly the same way as the deprecated **allow-configuration** statement at the same hierarchy level. [PR/720013: This issue has been resolved.]
- NTP-related **show** commands, such as **show ntp status**, might display incorrect output. [PR/722528: This issue has been resolved.]
- On EX4200 switches, after you issue the **request system zeroize media** command, you might not be able to establish a connection with the switch using SSH, and you might not be able to issue the **commit** command on the switch. [PR/723918: This issue has been resolved.]
- On EX8208 switches, when the Switch Fabric and Routing Engine (SRE) module is in the spare state and you configure it to go offline and then come back online again, the module's ST LED does not turn back on. [PR/724455: This issue has been resolved.]
- If you issue the **show krt next-hop** or **show krt iflist-next-hop** command, and if you later delete a route or the route is removed, an rpd core file might be created. [PR/727014: This issue has been resolved.]
- On EX8200 switches, after you issue the **request system zeroize media** command, the line cards might not come online. [PR/728082: This issue has been resolved.]
- If you include the **autoinstallation** configuration statement at the **[edit system]** hierarchy level, the switch interfaces might not work correctly. [PR/728344: This issue has been resolved.]
- The Ethernet switching process (eswd) might create a core file. [PR/732263: This issue has been resolved.]
- If you abruptly take a power supply offline, a chassis manager process (chassimd) core file might be created. [PR/737604: This issue has been resolved.]
- The **request system zeroize** command might not erase all files, such as files in the **/config**, **/var/db/config**, and **/var/db** directories. [PR/737916: This issue has been resolved.]
- If you use EZSetup to configure a root password that contains a comma (,), the characters after the comma are not checked during authentication, so it is possible to log in to the switch with several different passwords. As a workaround, configure the root password from the CLI. [PR/738310: This issue has been resolved.]
- When you quickly insert and then remove a line card, the chassis manager process (chassism) might become unstable. [PR/740730: This issue has been resolved.]

- On EX Series switches and SRX Series Services Gateways, when you enable "Change password every time the user logs out" on the active directory, the user is unable to change his or her password. [PR/740869: This issue has been resolved.]
- On EX8200 switches, a chassis manager process (chassism) core file might be created. [PR/745964: This issue has been resolved.]
- If you have configured PIM in dense or dense-sparse mode and there are more than 1500 sources for a group, a scheduler slip error (RPD\_SCHED\_SLIP) might occur, and IGMP might use a large number of CPU cycles. [PR/748420: This issue has been resolved.]
- When there is a large amount of NetBIOS traffic on the network, the switch might exhibit high latency while pinging between VLANs. [PR/748707: This issue has been resolved.]
- On EX4200 switches, a Packet Forwarding Engine process (pfem) core file might be created while the switch is running the PFE internal support script and saving the output to a file. [PR/749974: This issue has been resolved.]
- On all EX Series switches except EX8200 switches, if you have configured several policer settings in the same filter, they might all be overwritten when you change one of the settings. As a workaround, delete the setting and then add it back again with the desired changes. [PR/750497: This issue has been resolved.]
- You might see the following message in log files: **Kernel/ (COMPOSITE NEXT HOP) failed, err 6 (No Memory)**. [PR/751985: This issue has been resolved.]
- On EX3300 switches, if you configure more than 20 BGPv6 neighbor sessions, the command-line interface (CLI) might display the db> prompt. [PR/753261: This issue has been resolved.]
- On EX8200 switch line cards, a Packet Forwarding Engine process (pfem) core file might be created as the result of a memory segmentation fault. [PR/757108: This issue has been resolved.]
- On XRE200 External Routing Engines, when you issue the **show chassis hardware (<get-chassis-inventory>)** command, duplicate occurrences of <name> and <serial-number> tags under the <chassis> tag might result in malformed XML output. [PR/772507: This issue has been resolved.]
- When an EX Series switch is routing multicast traffic, that traffic might not exit from the multicast router port in the source VLAN. [PR/773787: This issue has been resolved.]
- EX4500 Series switches and EX8200-40XS line cards do not forward IP UDP packets when their destination port is 0x013f (PTP) or when the fragmented packet has the value 0x013f at the same offset (0x2c). [PR/775329: This issue has been resolved.]
- When EX Series switches receive packets across a GRE tunnel, the switches might not generate ARP packets. [PR/782323: This issue has been resolved.]
- After you remove an IPv6 interface configuration and then perform a rollback operation, the IPv4 label might change to explicit null. [PR/786537: This issue has been resolved.]

- When many packets are queued to have their next hop resolved, some packets might become corrupted. [PR/790201: This issue has been resolved.]
- If you configure IPv6 and VRRP, the IPv6 VRRP MAC address might be used incorrectly as the source MAC address when traffic is routed across VLANs. [PR/791586: This issue has been resolved.]

### **Interfaces**

- When you configure the **no-preempt** and **interface-tracking** options on a switch that is a VRRP master router, if the VRRP mastership is taken over by a switch that is a VRRP backup router and the tracking interface on the original master router goes down, then if the tracking interface on the original master router comes back up and the master's original priority is restored, the new master's mastership might transition to the original master router. [PR/699243: This issue has been resolved.]
- After multiple graceful Routing Engine switching (GRES) operations, the virtual management Ethernet (vme) interface might go down and then come up again after you issue the **restart ethernet-switching** command. [PR/719424: This issue has been resolved.]
- When you delete the VLAN mapping for an aggregated Ethernet (ae) interface, the Ethernet switching process (eswd) might crash and display the error message **No vlan matches vlan tag 116 for interface ae5.0**. [PR/731731: This issue has been resolved.]
- On EX8200 switches, the **master-only** configuration for the management interface does not work. [PR/753765: This issue has been resolved.]
- When EX Series switches receive packets across a GRE tunnel, they might not generate and send ARP packets to the device at the other end of the tunnel. [PR/782323: This issue has been resolved.]

### **J-Web Interface**

- In the J-Web interface on EX Series switches, J4350 Services Router, M Series routers, MX Series routers, and SRX210 Services Gateways, you cannot log out of the device using the CLI Terminal page (Diagnose > CLI Terminal), because the Logout option is not listed in the page. [PR/401772: This issue has been resolved.]
- If you have created dynamic VLANs by enabling MVRP from the CLI, then in the J-Web interface, the following features do not work with dynamic VLANs and static VLANs:
  - In the Port Configuration page (Configure > Interface > Ports)—Port profile (select the interface, click **Edit**, and select **Port Role**) or the VLAN option (select the interface, click **Edit**, and select **VLAN Options**).
  - VLAN option in the LACP (Link Aggregation Control Protocol) Configuration page (Configure > Interface > Link Aggregation)—Select the aggregated interface, click **Edit**, and click **VLAN**.
  - In the 802.1X Configuration page (Configure > Security > 802.1x)—VLAN assignment in the exclusion list (click **Exclusion List** and select **VLAN Assignment**) or the move to guest VLAN option (select the port, click **Edit**, select **802.1X Configuration**, and click the **Authentication** tab).

- Port security configuration page (Configure > Security > Port Security).
- In the Port Mirroring Configuration page (Configure > Security > Port Mirroring)—Analyzer VLAN or ingress or egress VLAN (click **Add** or **Edit** and then add or edit the VLAN).

[PR/669188: This issue has been resolved.]

- When a large number of inbound HTTP connections are established over an extended period of time, the HTTP process (httpd) might become trapped in a loop, resulting in high CPU utilization. The CPU load continues even after the stream of connection attempts is terminated. To reduce the CPU load, you must kill the process from the shell. Two workarounds are: disable the J-Web interface, or allow access to the J-Web interface only from trusted networks. Alternatively, apply a policer at the edge or on the control plane (lo0) to rate-limit inbound connections to TCP port 80. Note that the typical side effects of applying rate limiting to services (for example, an increased risk of successful DoS attacks) also apply to inbound J-Web interface connections, so care should be taken before making changes to control plane protection firewall filters. See RFC 6192 for guidance on protecting the router's control plane. [PR/693434: This issue has been resolved.]
- In the login/splash screen and the Help mapping file, the copyright date is set to 2011. [PR/731790: This issue has been resolved.]
- On the J-Web dashboard the Total number of ports field in the Capacity Utilization section might show incorrect values for a mixed EX4200 and EX4500 Virtual Chassis. As a workaround, use the **show chassis hardware | match PIC | except Virtual** command to display the correct values. [PR/734766: This issue has been resolved.]
- In the J-Web interface, if you click the EX8200-48T, EX8200-48F, or EX8200-8XS line card in the chassis view in the dashboard, the expanded line card might not load its interfaces and might not display the interface status for both the EX8208 and EX8216 switches. As a workaround, first click the EX8200-40XS in the same chassis view and then close that line card. Then, click the EX8200-48T, EX8200-48F, or EX8200-8XS line card to display the status of all interfaces. [PR/742448: This issue has been resolved.]
- If you used the CLI to create a redundant trunk link (RTG) group whose members are not trunk ports, you cannot edit this group from the J-Web interface. As a workaround, edit the group from the CLI. [PR/745458: This issue has been resolved.]
- For EX Series switches, when you use the J-Web interface software upload package, the **unlink** option does not work. [PR/746546: This issue has been resolved.]
- When a switch has no routed interfaces, you cannot use the J-Web interface to add OSPF areas. As a workaround, use the CLI to add these areas. [PR/746624: This issue has been resolved.]
- In the J-Web interface on an EX8200 switch that is set in virtual-chassis mode, when you expand the number of uplink modules, line cards that have no uplink module report an error or map ports to nonexistent modules. This problem happens the first time that you configure capacity utilization values. [PR/750854: This issue has been resolved.]

- The J-Web interface is vulnerable to HTML cross-site scripting attacks, also called XST or cross-site tracing. [PR/752398: This issue has been resolved.]
- When you configure the **no-tcp-reset** statement, the J-Web interface might be slow or unresponsive. [PR/754175: This issue has been resolved.]
- In the J-Web interface on EX Series switches and on M Series and MX Series routers, you might not be able to upload a configuration file from the Upload page (Maintain > Config Mgmt > Upload). [PR/784009: This issue has been resolved.]
- In the J-Web interface, the Help page for the Install package in the Software Maintenance page (Maintain > Software) might not appear. [PR/786654: This issue has been resolved.]

### ***Layer 2 and Layer 3 Protocols***

- On EX2200, EX3300, and EX6200 switches, and on EX8200 Virtual Chassis, NetBIOS snooping does not work. [PR/706588: This issue has been resolved.]
- If you try to configure a Layer 3 protocol such as IS-IS, OSPF, or RIP on a Layer 2 interface (that is, an interface configured with the family ethernet-switching), the commit operation fails. [PR/729923: This issue has been resolved.]
- When a BFD session has stale entries, it might flap. [PR/744302: This issue has been resolved.]

### ***Management and RMON***

- The connectivity-fault management (CFM) process (cfmd) might create a core file. [PR/597302: This issue has been resolved.]
- When you are using IS-IS for forwarding only IPv6 traffic and IPv4 routing is not configured, if you perform an SNMP get or walk on an IS-IS routing database table, the RPD process might crash and restart, possibly causing a momentary traffic drop. [PR/753936: This issue has been resolved.]
- When an SNMP string is longer than 30 characters, it is not displayed in Junos OS command output. [PR/781521: This issue has been resolved.]
- On Juniper Networks EX Series Ethernet Switches, M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers, after you upgrade to Junos OS Release 11.4R3, 11.4R4, or 12.1R2, the device might stop responding to SNMP ifIndex list queries. As a workaround, restart the device. If restarting the device is not an option, restart the shared-memory daemon (shm-rtssdbd). [PR/782231: This issue has been resolved.]

### ***MPLS***

- On EX3200 and EX4200 switches, no counters are incremented in MPLS statistics files for label-switched paths (LSPs) that are used for circuit cross-connects (CCCs). [PR/724371: This issue has been resolved.]

### ***Software Installation and Upgrade***

- When you use NSSU to upgrade from Junos OS Release 11.3R5, all traffic across a link aggregation group (LAG) might be dropped. [PR/733050: This issue has been resolved.]
- The **unlink** option in the **request system software add package unlink** command does not work on EX Series switches. [PR/739795: This issue has been resolved.]

### ***Unified Access Control (UAC)***

- When an EX Series switch is configured as a Junos OS enforcer on an IC Series Unified Access Control Appliance, the Odyssey Access Client (OAC) status might change from open/authenticating to open and authenticated. [PR/742369: This issue has been resolved.]

### ***Virtual Chassis***

- On EX4200 Virtual Chassis, if you delete an uplink interface on which the **family mpls** option is configured, the MPLS functionality on the corresponding symmetric interfaces on the other members might be affected. [PR/704480: This issue has been resolved.]
- On EX8200 Virtual Chassis, when you swap the members of a link aggregation group (LAG), a vmcore or ksyncd core file might be created on the backup Routing Engine. [PR/711679: This issue has been resolved.]
- In a setup in which two XRE200 External Routing Engines (one acting as the master, the other as the backup) are connected to a member of an EX8200 Virtual Chassis that has two Routing Engines (one acting as the master, the other as the backup), if you remove the master Routing Engine or if you reboot this Routing Engine (for example, using the **request system reboot member 0 re0** command when re0 is the master Routing Engine), interfaces on which the Link Aggregation Control Protocol (LACP) is configured might flap. This interface flapping does not occur if you remove or reboot the backup Routing Engine. [PR/718857: This issue has been resolved.]
- On EX4500 Virtual Chassis, you cannot configure a mastership priority of 0 from the J-Web interface. As a workaround, configure this priority from the CLI. [PR/721426: This issue has been resolved.]
- The XRE200 External Routing Engine temperature monitors, which you can view using the **show chassis environment** command, might report temperatures that are twice as high as the actual temperature. This temperature-reporting error has no impact on XRE200 External Routing Engine behavior. The fans and the system receive the correct temperature internally, so unwanted fan speed changes or an XRE200 External Routing Engine shutdown cannot occur as a result of this misreported temperature. However, the incorrect reported temperatures generate alarms and alarm messages. [PR/734233: This issue has been resolved.]



- In some help files, the copyright date is set to 2011 instead of 2012. [PR/735607: This issue has been resolved.]
- On EX4200 switches and EX4200 Virtual Chassis, the event process (eventd) might create a core file. [PR/737893: This issue has been resolved.]
- On EX3300 switches, when a Virtual Chassis is formed, the Virtual Chassis backup member's console CLI does not automatically display the Virtual Chassis master's console CLI. As a workaround, manually log out from the Virtual Chassis backup member. [PR/744241: This issue has been resolved.]
- When you configure EX4200 Virtual Chassis with automated installation scripts, the installation might fail. As a workaround, include the **member** option in the **request system scripts add** command. [PR/747476: This issue has been resolved.]
- On EX8200 standalone switches or EX8200 Virtual Chassis on which an aggregated Ethernet interface is configured, multiple core files might be created on the line cards. [PR/749298: This issue has been resolved.]
- In EX8200 Virtual Chassis, the switch might incorrectly send untagged packets. As a result, some hosts in the VLAN might experience connectivity issues. [PR/752021: This issue has been resolved.]
- In EX8200 Virtual Chassis, after one Virtual Chassis member is rebooted, the line card of the corresponding rebooted member switch is not brought down immediately, and hence the peer sees that the interfaces remain in the Up state. Additionally, the interface state is not cleared immediately in the switch card chassis kernel. The result is that the protocol session goes down, and traffic loss occurs even if you have configured nonstop active routing (NSR). [PR/754603: This issue has been resolved.]
- On XRE200 External Routing Engines, a chassis core file might be created. [PR/791959: This issue has been resolved.]

---

### Issues Resolved in Release 12.2R2

The following issues have been resolved since Junos OS Release 12.2R1. The identifier following the description is the tracking number in our bug database.

#### ***Access Control and Port Security***

- On EX6200 switches, LLDP stops working if you execute the **set ethernet-switching-options voip interface access-ports vlan** command. [PR/829898: This issue has been resolved.]

#### ***Class of Service***

- When you are configuring class-of-service (CoS) drop profiles, the commit operation might fail and might display the message **Missing mandatory statement: 'drop-probability'**. [PR/807885: This issue has been resolved.]

#### ***Ethernet Switching and Spanning Trees***

- You cannot configure a VLAN whose name contains a hyphen (-). As a workaround, use an underscore (\_) in the name instead. [PR/753090: This issue has been resolved.]

- Link-protection switchover and revertive mode might not work as expected. [PR/781493: This issue has been resolved.]
- Ethernet ring protection switching (ERPS; G.8032) does not block PVST BPDUs. [PR/793891: This issue has been resolved.]

### **Firewall Filters**

- If you apply a policer to an interface, the policer might not work, and messages similar to the following are logged: **dfw\_bind\_policer\_template\_to\_filter:205 Binding policer fails.** [PR/802489: This issue has been resolved.]
- On EX8200 Virtual Chassis, if you add and then delete a firewall filter for traffic that enters on one Virtual Chassis member and is transmitted out another member, IPv6 traffic might be dropped. If the ingress and egress interfaces are on the same member, the firewall filter works correctly. [PR/803845: This issue has been resolved.]

### **Hardware**

- On EX4550 switches, link autonegotiation does not work on 1-Gb SFP interfaces. [PR/795626: This issue has been resolved.]
- Non-Juniper Networks DAC cables do not work on EX Series switches. [PR/808139: This issue has been resolved.]
- The backlight on the LCD panel of EX4550 switches does not turn on. [PR/820473: This issue has been resolved.]

### **High Availability**

- After you perform a nonstop software upgrade (NSSU), you might notice a traffic outage of 150 seconds while the line cards are restarting. [PR/800460: This issue has been resolved.]

### **Infrastructure**

- The **wildcard range unprotect** configuration statement might not be synchronized with the backup Routing Engine. [PR/735221: This issue has been resolved.]
- After you successfully install Junos OS, if you uninstall AI scripts, an mgd core file might be created. [PR/740554: This issue has been resolved.]
- After an EX Series switch has been up for several days, the switch or FPC might reach 100 percent CPU usage and then stay at 100 percent. [PR/752454: This issue has been resolved.]
- The Junos OS kernel might crash because of a timing issue in the `ttymodem()` internal I/O processing routine. The crash can be triggered by simple remote access (such as Telnet or SSH) to the device. [PR/755448: This issue has been resolved.]
- While multicast is resolving routes, the following SPF-related error might be displayed: **SPF:spf\_change\_sre(),383: jt\_change () returned error-code(Not found:4)!** [PR/774675: This issue has been resolved.]
- When you issue the **show vrrp brief** command, a VRRP process (vrrpd) core file might be created. [PR/782227: This issue has been resolved.]

- On EX4550 switches, if you configure the management (me0) interface and a static route, the switch is unable to connect to a gateway. [PR/786184: This issue has been resolved.]
- When you add a new virtual routing and forwarding (VRF) instance, existing firewall filters might not be applied to the new VRF instance. [PR/786662: This issue has been resolved.]
- On XRE200 External Routing Engines on which DHCP snooping and dynamic ARP inspection are enabled, when packets are transmitted out a different line card type from the ingress interface, an SFID core file might be created. [PR/794293: This issue has been resolved.]
- On EX8200 switches, when you issue the **request system reboot other-routing-engine** command, a timeout error might be displayed before the Routing Engine initiates its reboot operation. [PR/795884: This issue has been resolved.]
- In MPLS implementations on EX Series switches, EXP bits that are exiting the provider edge switch are copied to the three least-significant bits of DSCP—that is, to IP precedence—rather than to the most-significant bits. [PR/799775: This issue has been resolved.]
- On EX3300 switches, when you are configuring BGP authentication, after you have configured the authentication key, BGP peering is never established. [PR/803929: This issue has been resolved.]
- EX4550 switches might not load the configuration file after you perform an automatic image upgrade. [PR/808964: This issue has been resolved.]
- On EX Series switches that have Power over Ethernet (PoE) capability, chassisd (the chassis daemon) might crash when running SNMP requests (for example, SNMP get, get-next, and walk requests) on pethMainPse objects. This is caused by the system trying to free memory that is already freed. As a workaround, avoid running SNMP requests on pethMainPse objects. [PR/817311: This issue has been resolved.]
- When an uplink module in the switch is operating in 1-gigabit mode, a chassism core file might be created if you remove an SFP transceiver from one of the module's interfaces. As the chassism process restarts, all traffic passing through the interface is dropped. This problem happens with both copper and fiber SFPs. [PR/828935: This issue has been resolved.]

### **Interfaces**

- EX4200 and EX4500 switches support 64 aggregated Ethernet interfaces even though the hardware can support 111 interfaces. [PR/746239: This issue has been resolved.]
- On EX Series switches, if you have configured a link aggregation group (LAG) with link protection, an interface on the backup member might drop ingress traffic. [PR/796348: This issue has been resolved.]
- An interface on an EX4550-32F switch might go up and down randomly even when no cable is plugged in. [PR/803578: This issue has been resolved.]
- On EX3300 switches, when you configure VRRP with MD5 authentication with the **preempt** option on a routed VLAN interface (RVI), a vmcore file might be created. As

a workaround, delete the **preempt** option and disable MD5 authentication for VRRP. [PR/808839: This issue has been resolved.]

### ***J-Web Interface***

- If a Virtual Chassis contains more than six members, the Support Information page (Maintain > Customer Support > Support information) might not load. [PR/777372: This issue has been resolved.]
- Some component names shown by the tooltip on the Temperature in the Health Status panel of the dashboard might be truncated. As a result, you might see many components that have the same name displayed. For example, the components **GEPHY Front Left**, **GEPHY Front Middle**, and **GEPHY Front Right** might all be displayed as **GEPHYFront**. [PR/778313: This issue has been resolved.]
- In a mixed EX4200 and EX4500 Virtual Chassis, the master chassis view might display the temperature indicator of the backup. [PR/783052: This issue has been resolved.]
- If you issue the **set protocols rstp interface logical-interface-name edge** configuration command from the command-line interface (CLI), the J-Web interface might show that the configuration in the Configuration detail for Desktop and Phone page is not applicable for the port profile. However, no functionality for the Desktop and Phone port profile is affected. [PR/791323: This issue has been resolved.]
- In the J-Web interface, if you enable a spanning-tree protocol (STP, RSTP, or MSTP) and then exclude some ports from the spanning tree, you might not be able to include these ports as part of a redundant trunk group (RTG). [PR/791759: This issue has been resolved.]
- In the J-Web interface on EX4500 and EX4550 switches, you can configure temporal and exact-temporal buffers, which are not supported by Junos OS. [PR/796719: This issue has been resolved.]
- In a mixed Virtual Chassis in which an EX4550 switch is the master and at least one Virtual Chassis member supports PoE, if you click **Configure > POE** and then click another tab, a javascript error might be displayed. [PR/797256: This issue has been resolved.]
- In the J-Web interface on EX4550 switches, if you are using in-band management and select EZSetup, the error message **undefined configuration delivery failed** is displayed even though the configuration has been successfully committed. [PR/800523: This issue has been resolved.]

### ***Layer 2 and Layer 3 Protocols***

- After an NSSU operation, OSPF might remain in the INIT state because the flooding entry is not programmed correctly. [PR/811178: This issue has been resolved.]

### ***Management and RMON***

- The incorrect ifType might be displayed for counters on physical interfaces. [PR/784620: This issue has been resolved.]
- After a Routing Engine switchover, LACP and MIB process (mib2d) core files might be created. [PR/790966: This issue has been resolved.]

- In EX3300 Virtual Chassis, if you perform an SNMP poll of jnxOperatingState for fan operation, the information for the last two members in the Virtual Chassis is incorrect. [PR/813881: This issue has been resolved.]

### ***Multicast Protocols***

- On XRE200 External Routing Engines on which PIM is configured, an NSSU operation might fail when performed when an MSDP peer is not yet up. As a workaround, either disable NSR for PIM using the **set protocols pim nonstop-routing disable** configuration comment or ensure that MSDP has reached the Established state before starting an NSSU operation. [PR/799137: This issue has been resolved.]

### ***Software Upgrade and Installation***

- After you upgrade Junos OS, a ppmmd core might be created, and protocols that use ppmmd might not work correctly. [PR/802315: This issue has been resolved.]

### ***Virtual Chassis***

- On XRE200 External Routing Engines, a chassism core file might be created. [PR/791959: This issue has been resolved.]

## **Issues Resolved in Release 12.2R3**

---

The following issues have been resolved since Junos OS Release 12.2R2. The identifier following the description is the tracking number in our bug database.

### ***Access Control and Port Security***

- Traffic leaks might occur for unknown unicast traffic and broadcast traffic from multiple VLANs when a MAC-RADIUS-assigned VLAN is set on a switch interface through a server-initiated attribute change. If the 802.1X interface has VLAN 100 assigned and the RADIUS server sends a different VLAN attribute (for example, 200 instead of 100), after the interface is assigned in VLAN 200, it also sends egress unknown unicast and broadcast traffic that belongs to VLAN 100. [PR/829436: This issue has been resolved.]

### ***Infrastructure***

- On EX8200 Virtual Chassis, when you swap the members of a LAG, a vmcore or ksyncd core file might be created on the backup Routing Engine. [PR/793778: This issue has been resolved.]
- On EX4200 switches, high CPU usage might be due to console cable noise. [PR/818157: This issue has been resolved.]
- On EX8200 Virtual Chassis, when both dscp and ieee-802.1 rewrite rules are applied on an RVI, deleting the filters and binding again on the same RVI or clearing interface statistics might create a pfem core file. [PR/828661: This issue has been resolved.]
- Multicast packets might be lost when the user switches from one IPv6 channel to another. [PR/835538: This issue has been resolved.]

- An SNMP poll might not return clear information for some field-replaceable units (FRUs), such as fans and power supplies. The FRU description might not indicate which physical switch contains the FRU. [PR/837322: This issue has been resolved.]
- If you reboot the switch with the RVI disabled, then even if you reenables the RVI, the RVI traffic is not routed in the Packet Forwarding Engine; the traffic is trapped to the CPU and is policed by the rate limit in the Packet Forwarding Engine. [PR/838581: This issue has been resolved.]

### ***Layer 2 and Layer 3 Protocols***

- On an EX4200 switch configured for VLAN translation, Windows NetBIOS traffic might not be translated. [PR/791131: This issue has been resolved.]

### ***Management and RMON***

- Some sFlow monitoring technology packets are dropped when the sFlow packet size exceeds 1500 bytes. This was due to a modification made to the sFlow packet size to align with the buffer descriptor size. [PR/813879: This issue has been resolved.]
- On EX8200 platforms, sFlow monitoring technology packets are generated with a bogus source MAC address of 20:0b:ca:fe:5f:10. The EX8200 platforms now use the outbound port's MAC address as the source MAC address for the sFlow monitoring technology traffic. [PR/815366: This issue has been resolved.]

## **Issues Resolved in Release 12.2R4**

---

The following issues have been resolved since Junos OS Release 12.2R3. The identifier following the description is the tracking number in our bug database.

### ***Access Control and Port Security***

- On EX Series switches, the LLDP-MED media endpoint class is shown as invalid. This problem is just a display issue—there is no functional impact. [PR/840915: This issue has been resolved.]

### ***Ethernet Switching and Spanning Trees***

- If an EX Series switch has a redundant trunk group (RTG) link, a MAC Refresh message might be sent on a new active link of the RTG when RTG failover occurs. The switch sends the RTG MAC Refresh message with a VLAN tag even though RTGs are configured on access ports. As a workaround, configure an RTG on trunk ports. [PR/853911: This issue has been resolved.]

### ***Hardware***

- On EX2200 switches, the RPS system is not recognized after you reseal the power supply, and configuration changes related to the RPS system result in an error message. [PR841785: This issue has been resolved.]
- On EX3200, EX4200, and EX8200 switches, the receiver signal average optical power is shown as 0.0000 in output for the **show interfaces diagnostics optics** command. The problem has been observed for SFP-SX and SFP-LX10 transceivers. [PR/854726: This issue has been resolved.]

### ***Infrastructure***

- The output of the **show system users no-resolve** command displays the resolved hostname. [PR/672599]
- For all EX Series switches except the EX8200 switch, hash index collisions are causing problems with the learning of MAC addresses in the forwarding database (FDB). You can now increase the maximum number of searchable hash indexes in increments of 4, from 4 to a maximum of 32 entries, using the CLI command **set ethernet-switching-options max-lookup-length**. [PR/842439: This issue has been resolved.]
- On EX8200 switches, the **commit synchronize** command fails and displays the error message: **error: could not open configuration database (juniper.data+)**. [PR/844315]
- In a mixed EX4200 and EX4500 Virtual Chassis, link aggregation might cause a PFEM core in some member switches. [PR/846498: This issue has been resolved.]
- When you boot up an EX2200 or EX3300 switch with Junos OS Release 12.2R1 or later, the message **?dog: ERROR - reset of uninitialized watchdog** appears. The message appears even if you reboot the switch by using the proper reboot procedure. The error does not cause a system reset; thus, you can ignore this message. [PR/847469: This issue has been resolved.]
- On EX Series switches, EXP CoS classification does not occur if EXP CoS classifiers are deleted and then added. [PR/848273: This issue has been resolved.]
- On EX4200 Virtual Chassis, any operation performed in private mode after the system is brought up with a scaled configuration creates an mgd core file. [PR/855990: This issue has been resolved.]

### ***Layer 2 and Layer 3 Protocols***

- On EX Series switches, the Cisco Discovery Protocol (CDP) and the VLAN Trunking Protocol (VTP) do not work through Layer 2 protocol tunneling (L2PT). [PR/842852: This issue has been resolved.]
- On EX Series switches, the Q-BRIDGE-MIB OID 1.3.6.1.2.1.17.7 reports the VLAN internal index instead of the VLAN ID. [PR/850299: This issue has been resolved.]

### ***Management and RMON***

- On EX Series switches, a configured OAM threshold value might be reset when the chassis is rebooted. [PR/829649: This issue has been resolved.]
- On EX4200 and EX4500 switches, adaptive sampling is triggered on interfaces configured for sFlow monitoring technology even though the sampling rate is less than 300 pps. [PR/840858: This issue has been resolved.]
- On EX4200 Virtual Chassis, the SNMP query or walk on ipNetToMediaPhysAddress does not match **show arp** command output. [PR/850051: This issue has been resolved.]

### ***Virtual Chassis***

- On EX Series switches, if you configure a physical interface's maximum transmission unit (MTU) with a large value and you do not reconfigure the family inet MTU, OSPF packets might be dropped when they reach the internal logical interface if the packet size exceeds 1900 bytes. All communications traffic between Routing Engines and between FPCs passes through the internal logical interface. The OSPF neighbor does not receive the OSPF transmissions and ends the OSPF session. The switch displays the error message **bmeb\_rx failed**. As a workaround, if possible, configure the family inet MTU of OSPF interfaces with no more than 1800 bytes, as in the following example: **set interfaces ge-0/0/1 unit 0 family inet mtu 1800**. [PR/843583: This issue has been resolved.]
- On EX4550 switches, dedicated Virtual Chassis ports (VCPs) are automatically assigned a new trunk ID when added to a link aggregation group (LAG). [PR/665876: This issue has been resolved.]

---

### **Issues Resolved in Release 12.2R5**

The following issues have been resolved since Junos OS Release 12.2R4. The identifier following the description is the tracking number in our bug database.

#### ***Access Control and Port Security***

- On EX4500, EX4550 and EX6200 switches, DHCPv6 is now supported. [PR/820403: This issue has been resolved.]
- On EX Series switches, when you use Zero Touch Provisioning, the switch is expected to reboot with the factory default configuration if there is no response from the DHCP server. Instead, the switch reverts to a configuration that was missing the **[edit system syslog]** stanza. [PR/857872: This issue has been resolved.]
- On EX Series switches, DHCP snooping binding does not renew the lease time when IPv6 is configured on the client VLAN. When DHCP snooping is configured with ARP inspection and when a client renews the lease, the switch does not update the DHCP snooping table with the new lease time. The lease eventually times out from the DHCP snooping table, and the client still has a valid lease. The client's ARP request eventually times out of the switch, and the client loses connectivity because ARP inspection blocks the ARP packet because of the client's missing entry in the DHCP snooping table. As a workaround, disable and then reenabale the client interface, or remove IPv6 from the VLAN. [PR/864078: This issue has been resolved.]

#### ***Ethernet Switching and Spanning Trees***

- On EX Series switches, when you issue the **show spanning-tree interface vlan-id vlan-id detail** command, the **vlan-id** parameter is ignored, and the output displays information for all interfaces instead of only for interfaces that are associated with the VLAN ID. [PR/853632: This issue has been resolved.]
- On EX Series switches, when a topology change is detected on an MSTP-enabled port, there might be a delay of several seconds before a BPDU is sent out with a topology change flag to all the other ports. When such a change is detected on an RSTP-enabled



port, a BPDU is sent out immediately with the topology change flag. [PR/860748: This issue has been resolved.]

### **Firewall Filters**

- On EX4200 switches, an aggregated Ethernet interface is not supported as a match condition in a firewall filter. [PR/886476: This issue has been resolved.]

### **Infrastructure**

- Rate limiting for management traffic (namely, SSH and Telnet) arriving on network ports causes file transfer speeds to be slow. [PR/831545: This issue has been resolved.]
- On EX8200 switches, multiple rpd process core files might be created on the backup Routing Engine after a nonstop software upgrade (NSSU) has been performed while multicast traffic is on the switch. [PR/841848: This issue has been resolved.]
- On EX4500 switches, the queue counters for child members are not updated while the **monitor interface ae** command is running. As a workaround, use the CLI command **monitor interface traffic**. [PR/846059]
- On EX4200 Virtual Chassis, a **/var partition is full** alarm and a **CHASSISD\_RE\_CONSOLE\_ME\_STORM** log might occur, caused by a console error storm, even though the **/var partition** is not full. You can ignore this alarm; it has no effect on the system. [PR/866863: This issue has been resolved.]
- On EX4500 switches and EX4500 Virtual Chassis, MPLS CoS classifications and rewrites might not work. [PR/869054: This issue has been resolved.]
- On EX2200 and EX3300 switches, storm control does not limit traffic to the set value when that traffic enters through uplink ports; the traffic is limited to 10 times the set value. [PR/879798: This issue has been resolved.]

### **Interfaces**

- Internal interfaces can sometimes fail autonegotiation after a reboot. The failed interfaces can be recovered using the **ifconfig up** command at the shell, but there is now an automatic detection and recovery mechanism to restore the interfaces if needed. [PR/829521: This issue has been resolved.]
- On EX3200 and EX4200 switches, high traffic on management Ethernet (me0) interfaces might affect switch control and management plane functions. [PR/876110: This issue has been resolved.]

### **High Availability**

- On EX8200 Virtual Chassis, an NSSU might fail. [PR/871288]
- On EX Series switches, there is no equivalent RPC available for the **show chassis nonstop-upgrade** command.

### **Layer 2 and Layer 3 Protocols**

- If you have configured NSR for PIM, a core file might be created on an upstream router because of high churn in unicast routes or a continuous clearing of PIM join-distribution

in the downstream router. To prevent this possibility, disable NSR for PIM. [PR/707900: This issue has been resolved.]

- On a device that is running PIM and with NSR enabled on the device, if a PIM corresponding interface flaps continuously, a PIM thread might attempt to free a pointer that has already been freed. This attempt causes the routing protocol daemon (rpd) to crash and create a core file. [PR/801104: This issue has been resolved.]
- On EX Series switches, if a multicast group address is configured with an invalid subnet, issuing the **commit** or **commit check** command creates a core file of the routing protocol daemon (rpd). This issue will result in the following error messages:

```
user@router#set routing-options multicast ssm-groups 224/1
```

```
[edit]
```

```
user@router# commit check
```

```
error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped  
core(0x86)
```

```
error: configuration check-out failed
```

```
[edit]
```

```
user@router# commit
```

```
error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped  
core(0x86)
```

```
error: configuration check-out failed
```

[PR/856925: This issue has been resolved.]

### ***Management and RMON***

- Under certain conditions, duplicate SNMP indexes might be assigned to different interfaces. This might cause the mib2d (Management Information Base II daemon) and the lacpd (LACP daemon) to crash and create core files. [PR/836823: This issue has been resolved.]
- On EX Series switches, when the ARP table is cleared from the CLI, the SNMP MIB ipNetToMediaPhysAddress might have more entries than the ARP table. [PR/853536: This issue has been resolved.]
- The sFlow monitoring technology feature is not supported on EX2200, EX2200-C, and EX3300 switches. [PR/872292: This issue has been resolved.]

### ***Multicast***

- On EX4500 switches, multicast packet fragments might be dropped. [PR/835855: This issue has been resolved.]

### ***Software Installation and Upgrade***

- On the EX2200-24T-DC-4G switch, autoinstallation is not activated during initial installation. Failure of autoinstallation impacts only the DC version EX2200-24T model. The failure is due to a missing configuration file for this model. [PR/873689: This issue has been resolved.]

### ***Virtual Chassis***

- The **request system scripts add** command does not install the AI-Scripts bundle package on all nodes of an EX8200 Virtual Chassis. [PR/832975: This issue has been resolved.]
- On EX4200 Virtual Chassis, if the MAC persistence timer is configured for 0 minutes, the system MAC base address is changed by the **request chassis routing-engine master switch** command when a master switchover occurs. The allowed values for mac-persistence-timer have been changed to a range of 1 to 60 minutes instead of 0 to 60 minutes. [PR/858330: This issue has been resolved.]
- On EX8200 Virtual Chassis, NetBIOS traffic might be dropped when it crosses the Virtual Chassis port extension (VCPe) connections. The NetBIOS traffic is dropped because of a conflict on the Packet Forwarding Engine of the Virtual Chassis member with the VCPe ports. [PR/877503: This issue has been resolved.]

---

## **Issues Resolved in Release 12.2R6**

The following issues have been resolved since Junos OS Release 12.2R5. The identifier following the description is the tracking number in our bug database.

### ***Access Control and Port Security***

- On an EX Series switch, when you configure LLDP-MED on a trunk interface and set that interface as a member of both a voice VLAN and another VLAN, and you then change the mode of that interface to port (access) mode, the switch might send two different voice VLAN TLVs in an LLDP advertisement, and a VoIP phone connected to that interface might randomly select a VLAN to join. Use the **monitor traffic interface interface-name** command to check this issue. [PR/884177: This issue has been resolved.]

### ***Infrastructure***

- On EX3200 switches, an SNMP trap for pethPsePortDetectionStatus is not sent when a VoIP phone is disconnected from a PoE port. [PR/877768: This issue has been resolved.]
- EX4200 switches do not form Virtual Chassis links over uplink ports that contain copper SFPs. [PR/881868: This issue has been resolved.]
- On EX Series switches, if you have configured a LAG with link protection, ingress traffic does not pass through the backup port. [PR/886205: This issue has been resolved.]

- On EX2200 switches, the CPU is completely consumed by swi7: clock and chassis processes when the Redundant Power System (RPS) is powered off but is connected to the switch. At the same time, link LEDs blink continuously. When the RPS is powered up, CPU utilization and switch function becomes normal. [PR/890194: This issue has been resolved.]
- EX4500 switches might reboot suddenly because they have accessed an invalid register value for a port; this problem might occur when you insert or remove SFPs, or exchange 10-gigabit and 1-gigabit SFPs in a specific port. [PR/891733: This issue has been resolved.]
- On EX Series switches, a primary file system corruption might not be detected and the system might not fail over to the backup partition. Some functional problems might occur. [PR/892089: This issue has been resolved.]
- On EX4200 switches, if you configure and apply more than 32 CoS rewrite rules, the Packet Forwarding Engine manager (pfem) creates core files continuously. [PR/893911: This issue has been resolved.]
- On EX8200 series switches equipped with EX8200-40XS line cards, when a port on the 40XS line card is connected to another device and the port is then disabled, the Carrier Transition count might increase continuously, which might cause high CPU utilization on EX8200 switches. [PR/898082: This issue has been resolved.]
- On EX4550 switches running Junos OS Release 12.2R5 or Release 12.3R3, any commit operations will cause a spike in CPU utilization. This might result in a timeout of LACP, BFD, and other protocols. [PR/898097: This issue has been resolved.]
- On EX2200 switches, the syslog message displays all IP addresses in reverse. For example, an ICMP packet from IP address 10.0.1.114 to 10.0.0.7 produces the syslog, **PFE\_FW\_SYSLOG\_IP: FW: ge-0/0/0.0 R icmp 114.1.0.10 7.0.0.10 0 0 (1 packets)**, but it should read **PFE\_FW\_SYSLOG\_IP: FW: ge-0/0/0.0 R icmp 10.0.1.114 10.0.0.7 0 0 (1 packets)**. [PR/898175: This issue has been resolved.]

### **Interfaces**

- On an EX3300 switch, when another vendor's AP is connected to one of the EX3300 interfaces, LLDP negotiation might fail and the AP is unable to boot. The system is storing the organization-specific TLV's OUI and subtype values in the parsed TLV-to-value buffer, and due to this, the offset for reading PoE power negotiation from the buffer has been changed.

As a workaround:

1. Unplug the AP.
2. Wait until the interface power goes to 0, and verify that the physical interface is down.
3. Issue the **set protocol lldp interface ap-interface-name power-negotiation disable** CLI command and commit the command.
4. Connect the AP.

The AP will power on in IEEE class mode (not negotiated power). [PR/898234: This issue has been resolved.]

- On EX Series switches, configuration of a static LACP system ID is not supported. [PR/889318: This issue has been resolved.]

#### ***Software Installation and Upgrade***

- When multiple EX8208 switches are upgraded to Junos OS Release 12.1R6, issuing the command **request system software delete jloader-ex-8200** will create a FIPS error core file. [PR/894987: This issue has been resolved.]

#### ***Virtual Chassis***

- On EX8200 Virtual Chassis, when NSB is enabled, continuously adding and deleting VLAN members along with continuously creating and deleting VLANs will cause a memory leak and create an eswd core file. [PR/878016: This issue has been resolved.]
- On EX Series Virtual Chassis, an upgrade with NSSU might cause a mismatch in the IFD index numbers between the master and backup PFEs. This will result in packets being dropped as they pass through the virtual chassis. [PR/882512: This issue has been resolved.]
- On EX8200 Virtual Chassis, during NSSU, all interfaces, including LAGs, might go down during FRU updates, resulting in traffic loss. [PR/893440: This issue has been resolved.]

#### **Issues Resolved in Release 12.2R7**

The following issues have been resolved since Junos OS Release 12.2R6. The identifier following the description is the tracking number in our bug database.

### ***Access Control and Port Security***

- On EX Series switches, if an interface enabled for 802.1x in single supplicant mode gets authenticated before being authorized by the RADIUS server, the dot1x process might crash and generate a core file. This issue can occur after an upgrade or mastership failover but is not seen every time. When this issue occurs, subsequent supplicant authentication of this interface will be affected. [PR/890536: This issue has been resolved.]

### ***Class of Service***

- On EX4200-48PX switch models, configuring the traffic shaping rate on an interface using the **set class-of-service interfaces *interface-name* shaping-rate** command might return the error message **shaping rate not allowed on interface *interface-name***. [PR/944172: This issue has been resolved.]

### ***Ethernet Switching and Spanning Trees***

- On EX Series switches, if RSTP is enabled and the EX Series switch is connected to third-party switches running STP, when a topology change notification (TCN) BPDU is received on the EX switch from those third-party switches, the EX Series switch will need to wait for 35 seconds (Forward Delay + Max Age) to acknowledge the TCN. As a result, the EX switch will keep receiving the TCN due to retransmission, and keep flushing the forwarding table during this period, causing traffic to be blocked. [PR/910136: This issue has been resolved.]

### ***Hardware***

- On EX2200 and EX3300 switches, for some types of SFP transceivers, the output of the **show interface diagnostics optics** CLI command contains an incorrect value of **0.0000 mW / - Inf dBm** for the **Receiver signal average optical power** field. [PR/909334: This issue has been resolved.]

### ***Infrastructure***

- On EX Series switches, a high rate of interface family (IFF) changes, such as access security related events, might cause a memory leak in the Software Forwarding Infrastructure Daemon (sfid), which is responsible for handling packets destined for the CPU. The memory leak might cause sfid process crashes with a core file generated. [PR/907968: This issue has been resolved.]
- On EX Series switches with Connectivity Fault Management (CFM) configured, if CFM is in distributed mode, CFM control packets are sent out on the best effort queue instead of the network control queue. When this issue occurs, the CFM packets could be dropped during periods of network congestion. [PR/923604: This issue has been resolved.]
- On an EX6200 switch, if you disconnect the master Routing Engine (RE0) and reconnect it, the backup Routing Engine (RE1) becomes the master, and then when the original RE0 is rebooted, it becomes the backup, but that new backup does not appear in the **show chassis routing-engine** command output on RE0 (the new master). [PR/919242: This issue has been resolved.]

- On EX Series switches, when the same filter is applied to multiple interfaces, if the filter is removed or deleted from one of the interfaces, the other filter might be destroyed, causing the Packet Forwarding Engine Manager to create a pfem process core file. [PR/927063: This issue has been resolved.]
- On EX Series switches that are running Junos OS Release 12.1R4 or 12.2R2 and higher, if you install AI-Scripts package releases earlier than 3.6R4 and 3.7R3 and then execute a reboot/commit sequence, the switch might generate a FIPS core file and will crash, resulting in login failure or interruption of service. [PR/940478: This issue has been resolved.]

### ***Software Installation and Upgrade***

- On EX Series switches, upgrading to Junos OS Release 12.2 might create an LLDP core file. [PR/809899: This issue has been resolved.]
- On EX8200 switches, an NSSU might cause some hosts to become unreachable because the ARP index for the impacted host route is incorrectly programmed. The host route references the old ARP index and fails to update the new ARP index. [PR/894436: This issue has been resolved.]

### ***Virtual Chassis***

- On EX4200 Virtual Chassis, adding a VLAN and configuring VSTP for the VLAN might create an eswd process core file. [PR/864100: This issue has been resolved.]
- On EX series Virtual Chassis, if a physically down Virtual Chassis port (VCP) is converted to a network port, broadcast and multicast traffic might be dropped on the VCP interface. [PR/905185: This issue has been resolved.]
- On EX4550 Virtual Chassis, SFPs might not be detected, causing continuous EEPROM read failed errors. [PR/911306: This issue has been resolved.]

### ***Issues Resolved in Release 12.2R8***

---

The following issues have been resolved since Junos OS Release 12.2R7. The identifier following the description is the tracking number in our bug database.

### ***Ethernet Switching and Spanning Trees***

- On EX Series switches except EX2200, when RSTP and VSTP are enabled at the same time, an RSTP topology change might delete MAC entries learned on VLANs managed by VSTP. [PR/900600: This issue has been resolved.]

### ***High Availability***

- On EX Series Virtual Chassis with a link aggregation group (LAG) configured, if one member link of the LAG is on the backup Routing Engine, traffic loss on the LAG interface might be observed during an NSSU. Traffic resumes after the GRES occurs in the last state of the NSSU. [PR/916352: This issue has been resolved.]

### ***Infrastructure***

- On EX4550 switches, high temperature alarms are not triggered based on the thresholds displayed in the output of the **show chassis temperature-thresholds** command, but on other internal thresholds. This problem affects only EX4550 switches and no other platforms. [PR/874506: This issue has been resolved.]
- On EX Series switches with a router firewall filter configured, the filter might not work if it is applied to an IPv6 VRRP-enabled interface; also, features corresponding to the filter, such as policers, do not work. [PR/926901: This issue has been resolved.]
- On an EX Series switch with TACACS+ authentication and accounting enabled, when the TACACS+ server is in an unresponsive state and sends an erroneous response with an End of File (EOF) that indicates that no data can be read from a data source, this circumstance causes the client to fail to decrement the sequence number that it manages locally. During that time, any TACACS+ authentication might fail. [PR/929273: This issue has been resolved.]
- When an SNMP walk is performed to query the native VLAN (mib-2.17.1.4.5.1...: dot1qPvid) or the logical type (trunk or access) of the interface (mib-2.17.1.4.3.1.5...: dot1qPortVlan), the SNMP walk might cause a memory leak on the Layer 2 address learning process (l2ald), and the process might crash with a core file generated. [PR/935981: This issue has been resolved.]
- On EX Series switches with filter-based forwarding configured, if there is any event that causes the Virtual Routing and Forwarding (VRF) route table to be deleted and added repeatedly, such as the VRF configuration being deleted and added repeatedly, this might cause the installation of routes to fail. [PR/949832: This issue has been resolved.]
- On EX2200, EX3200, EX3300, EX4200, EX4500, EX4550, and EX6200 switches, DHCPv6 unicast packets might get dropped after enabling a firewall filter on the loopback interface (lo0.0) to protect the Routing Engine. [PR/960687: This issue has been resolved.]



### *Virtual Chassis*

- In mixed EX4200 and EX4500 Virtual Chassis, the following log message appears after every reboot: **CHASSISD\_PIC\_OID\_UNKNOWN: Unable to find OID for PIC**. Also, for the PICs on the EX4200 switch, the jnxContentsType MIB might report the value jnxEX4500MediaCardSpacePIC.O. [PR/711871]

### Issues Resolved in Release 12.2R9

The following issues have been resolved since Junos OS Release 12.2R8. The identifier following the description is the tracking number in our bug database.

#### *Authentication and Access Control*

- On EX Series switches with 802.1X authentication enabled, when the RADIUS server is unreachable, the 802.1X-enabled interface might stop forwarding traffic if you deactivate the 802.1X protocol by deactivating the [edit protocols dot1x] stanza. As a workaround, deactivate the 802.1X protocol by issuing this command: **deactivate protocols dot1x authenticator interface**. [PR/947882: This issue has been resolved.]
- On an EX Series switch that has both 802.1X authentication (dot1x) and a dynamic firewall filter enabled, when the server-timeout value is set to a short time (for example, 3 seconds), if many clients try to authenticate at the same time, a **delay success authentication success** message might be received on the switch due to a RADIUS server timeout, the firewall filter might corrupt the interfaces on which the authentication attempts were made, and the subsequent client authentications might fail due to the stale firewall filter. As a workaround, configure a server-timeout value that is greater than 30 seconds. [PR/967922: This issue has been resolved.]

#### *Infrastructure*

- On EX Series switches that are configured for filter-based forwarding (FBF), FBF might not work as expected because the routing instance used in the FBF filter is not available in the Packet Forwarding Engine during bootup. To restore service, delete and then add the FBF filter again. [PR/952539: This issue has been resolved.]
- On EX Series switches, SNMP counters for aggregated Ethernet (ae) interfaces might not match the sum of the member links. [PR/957434: This issue has been resolved.]
- On an EX4500 or EX4550 switch with an MPLS circuit cross-connect (CCC) interface configured, there might be high CPU utilization by the software forwarding infrastructure daemon (sfid) while large amounts of IPv6 neighbor solicitation packets (for example, 1000 pps) are received on the MPLS CCC interface. [PR/961807: This issue has been resolved.]
- On EX8200 switches in a scaled MPLS scenario, the pfem process might create a core file because of conflicts in the memory table. [PR/967492: This issue has been resolved.]
- On EX8200 switches with CoS configured, when buffer-size temporal values are configured or changed multiple times while traffic on the wire is already oversubscribed, the corresponding queue might be locked and not forward traffic even after the traffic becomes normal.

As a workaround:

1. Power off the line card (FPC) where the temporal configuration is set.
2. Remove the temporal configuration and set the buffer percent.
3. Restart the FPC.

[PR/967984: This issue has been resolved.]

- On EX Series switches, the software forwarding infrastructure daemon (sfid) process might create a core file while processing a packet for which the TTL has expired, because the packet pointer is freed twice. [PR/988640: This issue has been resolved.]
- On EX Series switches running Junos OS Release 12.1R1 or later releases, the MPLS TTL might change to 1 on a transit MPLS switch, causing packets to be dropped on the egress MPLS tunnel due to TTL expiration. As a workaround, enable the **no-decrement-ttl** statement in the **[edit protocols mpls]** hierarchy. [PR/1005436: This issue has been resolved.]
- On EX Series switches, ARP reply packets might get dropped when the switch receives reverse-path forwarding (RPF) multicast failure packets at a high rate (for example, 300 pps). As a workaround, create a static ARP entry for the next-hop device. [PR/1007438: This issue has been resolved.]

#### ***Interfaces and Chassis***

- When you remove an SFP+ and then add it back or reboot the switch, and the corresponding disabled 10-gigabit interface is a member of a LAG, the link on that port might be activated.

As a workaround, when the port becomes active:

1. Delete the disable parameter from the interface and commit this configuration.
2. Disable the interface in the configuration and commit

[PR/947683: This issue has been resolved.]

### ***Layer 2 Features***

- On EX Series switches with L2PT and Q-in-Q tunneling enabled, MACs might not be learned for some addresses. The problem occurs when there is a high volume of L2PT packets. As a workaround, restart the eswd and sfid processes. [PR/996368: This issue has been resolved.]

### ***Layer 3 Features***

- On EX8200 switches, when the MTU value on a Layer 3 interface is configured as 1518 and you execute the **clear pim join** command or reboot the switch, multicast traffic might be dropped when packet sizes are greater than 1500, because the multicast route might eventually point to a smaller MTU value and packets cannot pass, even though the packet size is smaller than the MTU-configured value. As a workaround, configure all the Layer 3 interface MTUs to 9192. [PR/966704: This issue has been resolved.]

### ***Network Management and Monitoring***

- On EX Series switches, an OAM CFM interface might not recover automatically if the action in **[edit oam ethernet connectivity-fault-management action-profile link-down action action]** is **interface-down**. As a workaround, do not use **link-down** in the action profile. [PR/948082: This issue has been resolved.]

### ***Routing Protocols***

- On EX Series switches, the RPD process might create a core file due to PIM join/prune internal processing. [PR/817623: This issue has been resolved.]
- On EX Series switches, the routing protocol daemon (rpd) might generate a core file when multiple BGP sessions to neighbors in the same BGP peer group are forced to close status. Possible triggers are network reconvergence events causing the BGP sessions to go down, activating and deactivating the BGP protocol or the BGP peer-group configuration. [PR/823346: This issue has been resolved.]

### **Related Documentation**

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches on page 49](#)
- [Limitations in Junos OS Release 12.2 for EX Series Switches on page 50](#)
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches on page 57](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches on page 92](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches on page 94](#)

## Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches

- [Changes to Junos OS for EX Series Switches Documentation on page 92](#)
- [Errata on page 92](#)

### Changes to Junos OS for EX Series Switches Documentation

---

The following changes have been made to the documentation for Junos OS Release 12.2 for EX Series switches since it was published:

#### *Infrastructure*

- EX Series switches support IPv6 for the extended DHCP server and extended DHCP relay.
- The **show system services dhcp client** command on EX Series switches now includes decrementing-lease-time information to aid debugging. The field **Lease expires in:** has been added to the command output.

### Errata

---

This section lists outstanding issues with the published documentation for Junos OS Release 12.2 for EX Series switches.

- **auto-sw-update configuration statement**—The **auto-sw-update** configuration statement topic does not include information about the **ex-4200** and **ex-4500** options that were introduced for the statement in Junos OS Release 12.2. These options enable the automatic software update feature for all mixed Virtual Chassis that include EX4200 and EX4500 member switches.

You can use the instructions in [Configuring Automatic Software Update on Virtual Chassis Member Switches \(CLI Procedure\)](#) to enable the automatic software update feature for your mixed Virtual Chassis that includes EX4200 and EX4500 member switches.

The **auto-sw-update** configuration statement topic will be updated in a later release. [This issue is being tracked by PR/541092.]

- **Ethernet OAM link fault management**—You can configure Ethernet OAM link fault management (LFM) on aggregated interfaces.
- **Multicast load balancing on EX8200 switches**—On EX8200 switches, you can use the **show chassis multicast load-balance** command to see whether multicast load balancing is enabled, and if it is, what the hash mode has been set to. The command description will be added to the EX Series documentation in an upcoming release. [This issue was being tracked by PR/665072.]
- **request system software validate command**—The documentation for the **request system software validate** command incorrectly states that this command is supported on EX Series switches. This command is not supported on any EX Series switches. [This issue is being tracked by PR/803185.]
- **request system software add command**—The documentation for the **request system software add** command incorrectly states that the **validate** option is supported on EX

Series switches. This option is not supported on any EX Series switches. [This issue is being tracked by PR/821244.]

- **vllans configuration statement**—The documentation for the **vllans** configuration statement incorrectly states the required privilege levels as routing and routing-control. The correct privilege levels for this statement are system and system-control.
- The *Complete Hardware Guide for EX4550 Ethernet Switches* provides incorrect information about AC power cord specifications for an EX4550 switch. The correct specification is given in the topic [AC Power Cord Specifications for an EX4550 Switch](#).
- Zero Touch Provisioning (also known as EZ Touchless Provisioning) is introduced on EX3300 switches at Junos OS Release 12.2R5. The documentation incorrectly states that the feature was introduced on EX3300 switches at Junos OS Release 12.2R1. (The documentation will be updated to reflect the feature's name change from EZ Touchless Provisioning to Zero Touch Provisioning at an upcoming release.)
- The documentation for Junos OS Release 12.2 does not document the dedicated Virtual Chassis port link aggregation feature on EX4550 switches. The dedicated Virtual Chassis ports (VCPs) on EX4550 switches automatically form a link aggregation group (LAG) bundle when two or more dedicated VCPs are used to interconnect the same Virtual Chassis member switches starting in Junos OS Release 12.2R4. An EX4550 switch can include up to four dedicated VCPs, and all four dedicated VCPs can act as member links in a LAG when they are used to interconnect to the same Virtual Chassis member switch. Dedicated VCPs and optical ports configured as VCPs cannot be member links in the same LAG and are placed into different LAGs when both are configured to connect to the same EX4550 member switch.
- The *NETCONF XML Management Protocol Guide* incorrectly states that when performing a confirmed commit operation using the **<commit>** element, the **<confirm-timeout>** value specifies the number of minutes for the rollback deadline. The value of the **<confirm-timeout>** element actually specifies the number of seconds for the rollback deadline.

[*NETCONF XML Management Protocol Guide*]

#### Related Documentation

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches on page 49](#)
- [Limitations in Junos OS Release 12.2 for EX Series Switches on page 50](#)
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches on page 57](#)
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches on page 64](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches on page 94](#)

## Upgrade and Downgrade Instructions for Junos OS Release 12.2 for EX Series Switches

This section discusses the following topics:

- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 94](#)
- [Upgrading or Downgrading to Junos OS Release 12.2R1 on page 94](#)
- [Upgrading to Junos OS Release 12.1R2 or Later, with Existing VSTP Configurations on page 95](#)
- [Upgrading from Junos OS Release 10.4R3 or Later on page 95](#)
- [Upgrading from Junos OS Release 10.4R2 or Earlier on page 96](#)
- [Upgrading EX Series Switches Using NSSU on page 97](#)

---

### Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

For more information about EEOL releases and to review a list of EEOL releases, see [Junos Software Dates & Milestones](#).

---

### Upgrading or Downgrading to Junos OS Release 12.2R1

Before downgrading to Junos OS Release 12.2R1 from Release 12.2R2 or a later release, check your configuration to determine whether it contains the **interface all** statement or a family ethernet-switching interface for any Layer 3 routing protocol (at the **[edit protocols]** configuration hierarchy level). If these elements are present, you must delete them before performing the downgrade, so that errors do not occur during the downgrade operation.

Likewise, before upgrading to Junos OS Release 12.2R1 from an earlier release, make the same changes to your configuration file.

### Upgrading to Junos OS Release 12.1R2 or Later, with Existing VSTP Configurations

If you are upgrading to Junos OS Release 12.1R2 or later from Release 12.1R1 or earlier, ensure that any VSTP configurations on the switch meet the following guidelines. If the VSTP configurations do not meet these guidelines and you run the upgrade, the upgrade fails and you have to connect the console, change the invalid VSTP configurations, and commit the changed configurations through the console. Guidelines for VSTP configurations are:

- If you have specified physical interfaces for VSTP-configured VLANs, ensure that those interfaces are members of the VLANs specified in the VSTP configuration. If the VSTP configuration specifies `vlan all`, then the interfaces configured under `vstp vlan all` must be members of all VLANs.
- If the interfaces are not members of the VLANs in the VSTP configurations but are already added to the VSTP configurations, remove them from those configurations, add them to the VLANs, and then add them back to the VSTP configurations.

This issue is being tracked by PR/736488 in our bug database.

### Upgrading from Junos OS Release 10.4R3 or Later

This section contains the procedure for upgrading from Junos OS Release 10.4R3 or later to Junos OS Release 12.2. You can use this procedure to upgrade Junos OS on a standalone EX Series switch with a single Routing Engine and to upgrade all members of a Virtual Chassis or a single member of a Virtual Chassis.

To upgrade Junos OS on an EX6200 or EX8200 switch with dual Routing Engines, see [Installing Software on an EX Series Switch with Redundant Routing Engines \(CLI Procedure\)](#).

On switches with dual Routing Engines or on Virtual Chassis, you might also be able to use nonstop software upgrade (NSSU) to upgrade Junos OS. See [“Upgrading EX Series Switches Using NSSU” on page 97](#) for more information.

To upgrade Junos OS on a switch with a single Routing Engine or on a Virtual Chassis:

1. Download the software package as described in [Downloading Software Packages from Juniper Networks](#).
2. (Optional) Back up the current software configuration to a second storage option. See the [Junos OS Installation and Upgrade Guide](#) for instructions.
3. (Optional) Copy the software package to the switch. We recommend that you use FTP to copy the file to the `/var/tmp` directory.

This step is optional because you can also upgrade Junos OS using a software image that is stored at a remote location.

4. Install the new software package on the switch:

```
user@switch> request system software add package
```

Replace *package* with one of the following paths:

- `/var/tmp/package.tgz`—For a software package in a local directory on the switch

- `ftp://hostname/pathname/package.tgz` or  
`http://hostname/pathname/package.tgz`—For a software package on a remote server

`package.tgz` is the name of the package; for example,  
`jinstall-ex-4200-11.4R1.8-domestic-signed.tgz`.

To install software packages on all switches in a mixed EX4200 and EX4500 Virtual Chassis, use the **set** option to specify both the EX4200 package and the EX4500 package:

```
user@switch> request system software add set [package package]
```

To install the software package on only one member of a Virtual Chassis, include the **member** option:

```
user@switch> request system software add package member member-id
```

Other members of the Virtual Chassis are not affected. To install the software on all members of the Virtual Chassis, do not include the **member** option.



**NOTE:** To abort the installation, do not reboot your device. Instead, finish the installation, and then issue the `request system software delete package.tgz` command, where `package.tgz` is the name of the package; for example, `jinstall-ex-8200-11.4R1.8-domestic-signed.tgz`. This is the last chance to stop the installation.

5. Reboot the switch to start the new software:

```
user@switch> request system reboot
```

To reboot only a single member in a Virtual Chassis, include the **member** option:

```
user@switch> request system reboot member
```

6. After the reboot has finished, log in and verify that the new version of the software is properly installed:

```
user@switch> show version
```

7. Once you have verified that the new Junos OS version is working properly, copy the version to the alternate slice to ensure that if the system automatically boots from the backup partition, it uses the same Junos OS version:

```
user@switch> request system snapshot slice alternate
```

To update the alternate root partitions on all members of a Virtual Chassis, include the **all-members** option:

```
user@switch> request system snapshot slice alternate all-members
```

---

### Upgrading from Junos OS Release 10.4R2 or Earlier

To upgrade to Junos OS Release 12.2 from Release 10.4R2 or earlier, first upgrade to Junos OS Release 11.4 by following the instructions in the Junos OS Release 11.4 release notes. See *Upgrading from Junos OS Release 10.4R2 or Earlier* or *Upgrading from Junos OS Release 10.4R3 or Later* in the [Junos OS 11.4 Release Notes](#).



## Upgrading EX Series Switches Using NSSU

You can use NSSU to upgrade Junos OS releases on standalone EX6200 and EX8200 switches with dual Routing Engines and on EX3300, EX4200, EX4500, and EX8200 Virtual Chassis. For instructions on how to perform an upgrade using NSSU, see:

- [Upgrading Software on an EX3300 Virtual Chassis, EX4200 Virtual Chassis, EX4500 Virtual Chassis, or Mixed EX4200 and EX4500 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\)](#)
- [Upgrading Software on an EX6200 or EX8200 Standalone Switch Using Nonstop Software Upgrade \(CLI Procedure\)](#)
- [Upgrading Software on an EX8200 Virtual Chassis Using Nonstop Software Upgrade \(CLI Procedure\)](#)

Table 1 on page 97 details the switch platforms on which NSSU is supported and the required Junos OS releases.

**Table 1: Platform and Junos OS Upgrade Support for NSSU**

Switch Platform	Upgrade from Junos OS Release x.x	Upgrade to Junos OS Release 12.2
EX3300 Virtual Chassis	Releases earlier than 12.1R2	Not supported
	12.1R2 or later	Supported
EX4200 Virtual Chassis, EX4500 Virtual Chassis, and mixed EX4200 and EX4500 Virtual Chassis	Releases earlier than 12.1R1	Not supported
	12.1R1 or later	Supported
EX6200 standalone switch	Releases earlier than 12.1R2	Not supported
	12.1R2 or later	Supported
EX8200 standalone switch	10.4R1 or later	Not supported
	11.1R1 or later	Supported
	11.2R1 or later	Supported
	11.3R1 or later	Supported
	11.4R1 or later	Supported
	12.R1 or later	Supported

Table 1: Platform and Junos OS Upgrade Support for NSSU (*continued*)

Switch Platform	Upgrade from Junos OS Release x.x	Upgrade to Junos OS Release 12.2
EX8200 Virtual Chassis	10.4R1 or later	Not supported
	11.1R1, 11.1R2, or 11.1R3	Not recommended
	11.1R4 or later	Supported
	11.2R1 or later	Supported
	11.3R1 or later	Supported
	11.4R1 or later	Supported
	12.1R1 or later	Supported



**NOTE:** NSSU on an EX8200 Virtual Chassis is not recommended for Junos OS Release 12.2R5.

On an EX8200 Virtual Chassis, an NSSU operation can be performed only if you have configured the XRE200 External Routing Engine member ID to be 8 or 9.



**NOTE:** If you are using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 11.1 and sFlow technology is enabled, disable sFlow technology before you perform the upgrade using NSSU. After the upgrade is complete, you can reenables sFlow technology. If you do not disable sFlow technology before you perform the upgrade with NSSU, sFlow technology does not work properly. This issue does not affect upgrades from Junos OS Release 11.2 or later.



**NOTE:** If you are using NSSU to upgrade the software on an EX8200 switch from Junos OS Release 11.1 and NetBIOS snooping is enabled, disable NetBIOS snooping before you perform the upgrade using NSSU. After the upgrade is complete, you can reenables NetBIOS snooping. If you do not disable NetBIOS snooping before you perform the upgrade with NSSU, NetBIOS snooping does not work properly. This issue does not affect upgrades from Junos OS Release 11.2 or later.

#### Related Documentation

- [New Features in Junos OS Release 12.2 for EX Series Switches on page 41](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for EX Series Switches on page 49](#)

- [Limitations in Junos OS Release 12.2 for EX Series Switches on page 50](#)
- [Outstanding Issues in Junos OS Release 12.2 for EX Series Switches on page 57](#)
- [Resolved Issues in Junos OS Release 12.2 for EX Series Switches on page 64](#)
- [Changes to and Errata in Documentation for Junos OS Release 12.2 for EX Series Switches on page 92](#)

## Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers

---



**CAUTION:** This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for M, MX, and T Series routers. If you use unified ISSU when upgrading between releases, see the outstanding issues for [High Availability \(HA\) and Resiliency on page 214](#) in the current software release.

---

- [New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 100](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 154](#)
- [Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 170](#)
- [Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 171](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 285](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 322](#)

### New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers

The following features have been added to Junos OS Release 12.2. Following the description is the title of the manual or manuals to consult for further information.

- [Class of Service \(CoS\) on page 101](#)
- [Forwarding and Sampling on page 107](#)
- [High Availability \(HA\) and Resiliency on page 108](#)
- [Interfaces and Chassis on page 110](#)
- [Junos OS Installation and Upgrade on page 127](#)
- [Junos OS XML API and Scripting on page 128](#)
- [Layer 2 Ethernet Services on page 129](#)
- [Multiprotocol Label Switching \(MPLS\) on page 133](#)
- [Multicast on page 134](#)
- [Network Management and Monitoring on page 137](#)
- [Routing Policy and Firewall Filters on page 138](#)
- [Routing Protocols on page 140](#)
- [Services Applications on page 140](#)
- [Subscriber Access Management on page 140](#)

- [User Interface and Configuration on page 150](#)
- [VPNs on page 150](#)

### **Class of Service (CoS)**

---

- **Class-of-service features on the SONET/SDH OC192/STM64 MICs (MX Series routers)**—The following class-of-service (CoS) features are supported on the 1-port SONET/SDH OC192/STM64 MICs (model number: MIC-3D-IOC192-XFP):
  - The preclassifier block classifies the arriving packets. Packets are preclassified into traffic classes (drop precedence and scheduling priority of a packet depend on the traffic class) based on the sources of priority information.
  - Ingress behavior aggregate (BA) classification for DiffServ code point (DSCP), IP precedence, and MPLS EXP bits.
  - Shaping rates at the queue level.
  - Configurable bandwidth profiles with percentages.
  - Dynamic bandwidth allocation among different services.
  - Scheduler node scaling.
  - By default, eight egress queues are created on the physical interface. If per-unit scheduling is not configured, the same eight queues are shared across all logical interfaces.
  - Simple ingress policers.
  - On MPCs that do not support rich queuing, only coarse-grained queuing is provided.
  - On MPCs that do not support rich queuing, scheduling is available only at the physical interface level and not at the logical interface level. The per-unit scheduler cannot be configured on any of the physical interfaces on a MIC.
  - On MPCs that support rich queuing (for example, MX-MPC1-3D-Q), the **per-unit scheduler** statement can be configured on a physical interface only if the encapsulation is Frame Relay. If per-unit scheduling is configured, then each logical interface has eight queues.
  - Three levels of scheduling are supported on MPCs that support rich queuing:
    - Layer 1: Port level
    - Layer 2: Logical interface level
    - Layer 3: Queue level
  - Delay buffer allocation. By default, 100 ms worth of buffer is available on all the MPCs with or without rich queuing.
  - Parameters at the logical and physical interface levels: **guaranteed-rate**, **shaping-rate**, and **weighted-rate**.

[[BA Classifier Overview](#), [Scheduler Node Scaling on Trio MPC/MIC Interfaces Overview](#), [CoS on Trio MPC/MIC Features Overview](#), [per-unit-scheduler](#), [Providing a Guaranteed Minimum Rate](#), [shaping-rate](#)]

- **Class-of-service features supported on the T4000 Core Router**—The following class-of-service (CoS) features are supported on the T4000 Core Router (Type 5 FPCs):

Layer 3 Rewrite:

- IPv4 DSCP rewrite
- IPv4 INET Precedence rewrite
- IPv6 DSCP rewrite
- MPLS EXP rewrite
- Simultaneous MPLS EXP and IPv4 Precedence rewrite

In the case of L3VPN/L2VPN/VPLS, the following rules apply to simultaneous MPLS EXP and IPv4 precedence rewrite operation, under the **[edit class-of-service interfaces interface-name unit logical-unit-number rewrite-rules exp rewrite-rule-name] protocol protocol-types**; hierarchy:

The **protocol** statement defines the types of MPLS packets and possible packet configurations for the following options:

- mpls
- mpls-inet-both
- mpls-inet-both-non-vpn



**NOTE:** For L3VPN/L2VPN/VPLS, mpls-inet-both is not supported on the T4000.

---

Aggregated Ethernet:

- All CoS mechanisms that are supported on regular interfaces are supported on bundles.
- CoS with member links can be on different packet forwarding engines and line cards.

Shaping and Scheduling:

- Physical interface scheduling (eight queues per port)
  - Four packet loss priority levels
  - Unused bandwidth sharing among queues
  - Per physical interface shaping
- **Support for physical interface shaping on T4000 Routers (Type 5 FPCs)**—Enables a physical interface to shape traffic based on the rate-limited bandwidth of the total

interface bandwidth. This feature applies to physical interfaces on T4000 routers with Type 5 FPCs.

You can shape the output of a physical interface so that the interface transmits less traffic than it is physically capable of carrying.

To shape traffic on the physical interface, include the **shaping-rate** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level or include the **output-traffic-control-profile** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level.

[\[Applying a Shaping Rate to Physical Interfaces on T4000 Routers with Type 5 FPCs Overview\]](#)

- **CLI support for global and physical interface level classifiers**—Enables classification and rewrite at physical interface and global bind-points.

You can define EXP classification at a global level.

You can define the following features at the physical interface level:

- DSCP and inet-precedence classifiers
- DSCP and inet- precedence rewrites
- ieee-802.1 classifiers (inner and outer)
- ieee-802.1 rewrites (outer)

You can define the following features at the logical interface level:

- Fixed classification
- EXP rewrites

To configure global EXP classifiers, include the **classifiers exp *classifier-name*** statement at the **[edit class-of-service system-defaults]** hierarchy level.

To configure classifiers or rewrite rules at the physical interface, include either the **classifiers** or the **rewrite-rules** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level. To display classifiers configured under **system-defaults**, enter the **show class-of-service system-defaults** command.

To display classifiers and rewrites bound to physical interfaces, enter the **show class-of-service interfaces *Interface-name*** command.

[\[Classifiers and Rewrite Rules at the Global and Physical Interface Levels Overview\]](#)

- **Support for rate limiting on T4000 Type 5 FPC (T4000-FPC5-3D)**—You can configure rate limiting on T4000 routers with Type 5 FPC at the **[edit class-of-service schedulers *scheduler-name*]** hierarchy level.

[\[transmit-rate\]](#)

- **Support for hierarchical schedulers on aggregated Ethernet interfaces (MX Series routers)**—Enables you to apply hierarchical schedulers on aggregated Ethernet (AE) bundles through the use of interface-sets. This feature is supported at egress only on MX Series routers.

You can configure interface sets for AE interfaces created under static configurations.

You can configure class-of-service parameters on aggregated interfaces, in either link-protect or non-link-protect mode at the physical interface level. The CoS configuration is fully replicated for all AE member links in link-protect mode.

You can control the way these parameters are applied to member-links in non-link-protect mode by configuring the aggregated interface to operate in scaled mode or replicate mode. By default, scaled mode is used.

The link membership list and scheduler mode of the interface set is inherited from the underlying aggregated Ethernet interface over which the interface set is configured. When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is set to **member-link-scheduler replicate**, the scheduling parameters of the interface set are copied to each of the member links.

If the scheduler mode of the aggregated Ethernet interface is set to **member-link-scheduler scale**, the scheduling parameters are scaled based on the number of active member links (scaling factor is  $1/A$ , where  $A$  is the number of active links in the bundle) and applied to each of the aggregated interface member links.

To configure an interface set, include the **interface-set *interface-set-name*** statement at the **[edit class-of-service interfaces]** hierarchy level.

To apply scheduling and queuing parameters to the interface-set, include the **output-traffic-control-profile *profile-name*** statement at the **[edit class-of-service interfaces interface-set *interface-set-name*]** hierarchy level.

To apply the traffic control profile to the interface or interface set, include the **output-traffic-control-profile-remaining *profile-name*** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level, or the **[edit class-of-service interfaces *interface-name* interface-set *interface-set-name*]** hierarchy level, respectively.

[\[Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview\]](#)

- **Enhancement to the intelligent oversubscription feature on SONET/SDH OC48/STM16 IQE PICs**—Beginning with Junos OS Release 12.2, the support for maximum bandwidth optimization on SONET/SDH OC48/STM16 IQE PICs is increased to 300 percent with an additional priority group being created for all queues marked with low priority. When the sum of transmission rate for all queues exceeds 100 percent, the interface is in an oversubscribed state. At the time of oversubscription, the queues are split into three priority groups with the intelligent oversubscription feature enhancement:

- Strict High
- High, Medium-High, and Medium-Low
- Low

The sum of transmission rates for all queues in each of the previous priority groups is less than or equal to 100 percent, thereby allowing the SONET/SDH OC48/STM16 IQE PICs to support the maximum bandwidth optimization by overconfiguring the available bandwidth up to 300 percent.



- **Classification and DSCP marking of distributed protocol handler traffic**—The scope of traffic affected by the **host-outbound-traffic** statement is expanded. When it was introduced in Junos OS Release 8.4, the **host-outbound-traffic** statement at the **[edit class-of-service]** hierarchy level enabled you to specify the forwarding class assignment and DiffServ code point (DSCP) value for egress traffic sent from the Routing Engine. Affected traffic included control plane packets (such as OSPF hello and ICMP echo reply [ping] packets) and TCP-related packets (such as BGP and LDP control packets).

In Junos OS Release 12.2R2, the same configuration applies to *distributed protocol handler traffic* in addition to Routing Engine traffic. Distributed protocol handler traffic refers to traffic from the router's periodic packet management process (ppm) sessions, and it includes both IP (Layer 3) traffic such as BFD keepalive messages and non-IP (Layer 2) traffic such as LACP control traffic on aggregated Ethernet. DSCP changes do not apply to MPLS EXP bits or IEEE 802.1p bits. The specified queue must be correctly configured. The affected traffic includes distributed protocol handler traffic as well as Routing Engine traffic for egress interfaces hosted on MX Series routers with Packet Forwarding Engines, and on M120, M320, and T Series routers.

If you need the Routing Engine traffic and distributed protocol handler traffic to be classified in different forwarding classes or marked with different DSCP values, then you need to configure some additional steps. Apply a standard firewall filter to the loopback interface and configure the filter actions to set the forwarding class and DSCP values that override the **host-outbound-traffic** settings.

For interfaces on MX80 routers, LACP control traffic is sent through the Routing Engine rather than through the Packet Forwarding Engine.



**NOTE:** Any DSCP rewrite rules configured on a 10-Gigabit Ethernet LAN/WAN PIC with SFP+ overwrite the DSCP value rewritten as specified under the **host-outbound-traffic** statement.

The following partial configuration example classifies egress traffic from the Routing Engine as well as distributed protocol handler traffic:

```
[edit]
class-of-service {
  host-outbound-traffic {
    forwarding-class my_fc_control-traffic_dph;
    dscp-code-point 001010;
  }
  forwarding-classes {
    queue 5 my_fc_control-traffic_dph;
    queue 6 my_fc_control_traffic_re;
  }
}
interfaces {
  lo0 {
    unit 0 {
      family inet {
        filter {
          output my_filter_reclassify_re;
        }
      }
    }
  }
}
```

```
    }  
  }  
}  
firewall {  
  filter my_filter_reclassify_re {  
    term 1 {  
      then {  
        forwarding-class my_fc_control_traffic_re;  
        dscp code-points 000011;  
        accept;  
      }  
    }  
  }  
}
```

The statements in the example configuration cause the router to classify egress traffic from the Routing Engine and distributed protocol handler traffic as follows:

- Distributed protocol handler traffic is classified to the `my_fc_control-traffic_dph` forwarding class, which is mapped to queue 5. Of those packets, Layer 3 packets are marked at egress with DSCP bits 001010 (10 decimal), which is compatible with ToS bits 00101000 (40 decimal).
- Routing Engine traffic is classified to the `my_fc_control-traffic_re` forwarding class, which is mapped to queue 6. Of those packets, Layer 3 packets are marked at egress with DSCP bits 001100 (12 decimal), which is compatible with ToS bits 00110000 (48 decimal).

If you do not apply the firewall filter to the loopback interface, Routing Engine-sourced traffic is classified and marked using the forwarding class and DSCP value specified in the **host-outbound-traffic** configuration statement.

If you omit both the firewall filter and the **host-outbound-traffic** configuration shown in the previous configuration, then all network control traffic—including Routing Engine-sourced and distributed protocol handler traffic—uses output queue 3 (the default output queue for control traffic), and DSCP bits for Layer 3 packets are set to the default value 0 (Best Effort service).

#### [[Junos Class of Service Configuration](#)]

- **Enhancements to scheduler configuration on FRF.16 physical interfaces**—Starting with Release 12.2R2, Junos OS extends the class-of-service scheduler support on FRF.16 physical interfaces to the **excess-rate**, **excess-priority**, and **drop-profile-map** configurations. The **excess-rate**, **excess-priority**, and **drop-profile-map** statements are configured at the `[edit class-of-service schedulers scheduler-name]` hierarchy level.
  - Support for the **drop-profile-map** configuration enables you to configure random early detection (RED) on FRF.16 bundle physical interfaces.
  - Support for the **excess-rate** configuration enables you to specify the percentage of the excess bandwidth traffic to share.
  - Support for the **excess-priority** configuration enables you to specify the priority for excess bandwidth traffic on a scheduler.

This feature is supported only on multiservices PICs installed on MX Series routers.

- **Accurate reporting of output counters for MLFR UNI NNI bundles**—Starting with Release 12.2R2, Junos OS reports the actual output counters in the multilink frame relay (MLFR) UNI NNI bundle statistics section of the **show interfaces lsq-interface statistics** command output. From this release on, Junos OS also provides per-DLCI counters for logical interfaces. In earlier releases, there was a discrepancy between the actual output counters and the reported value because of errors in calculating the output counters at the logical interface level. That is, at the logical interface level, the output counter was calculated as the sum of frames egressing at the member links instead of providing the output counter as the sum of per-DLCI output frames

### Forwarding and Sampling

- **Host fast reroute**—Adds a precomputed protection path into the Packet Forwarding Engine, such that if a link between a provider edge device and a server farm becomes unusable for forwarding, the Packet Forwarding Engine can use another path without having to wait for the router or the protocols to provide updated forwarding information. Host fast reroute is a technology that protects IP endpoints on multipoint interfaces, such as Ethernet. This technology is important in data centers where fast service restoration for server endpoints is critical. After an interface or a link goes down, host fast reroute enables the local repair time to be approximately 50 milliseconds. You can configure Host fast reroute by adding the **link-protection** statement to the interface configuration in the routing instance. We recommend that you include this statement on all provider edge (PE) devices that are connected to server farms through multipoint interfaces.

[ *Example: Configuring Host Fast Reroute* ]

- **Distributed keepalive support from Packet Forwarding Engine to LNS PPP tunneled sessions on MPCs (MX Series routers)**—Junos OS supports client-initiated and server-initiated Point-to-Point Link Control Protocol (PPP LCP) echo request and reply packet handling from the Packet Forwarding Engine to the L2TP Network Server (LNS) PPP tunneled sessions on an MPC.

Keepalive aging timeout is defined as the product of the keepalive interval and down-count values at the LNS. If the keepalive aging timeout is greater than 180 seconds, the keepalive packets are handled at the Routing Engine. If the aging timeout is less than or equal to 180 seconds, the packets are handled at the Packet Forwarding Engine.



**NOTE:** When you scale the network to handle thousands of sessions, we recommend that you configure the keepalive aging timeout to be less than 180 seconds.

The display of the **show ppp interface extensive** command is enhanced to show the keepalive statistics for keepalives that are handled at the Routing Engine.

[ *show ppp interface* ]

- **Limiting traffic black-hole time on M320 routers by detecting Packet Forwarding Engine destinations that are unreachable over the fabric**—Enables M320 routers to limit traffic black-hole time by detecting unreachable destination Packet Forwarding Engines. The router signals neighboring routers when it cannot carry traffic because of the inability of some or all source Packet Forwarding Engines to forward traffic to some or all destination Packet Forwarding Engines on any fabric plane, after interfaces have been created. This inability to forward traffic results in a traffic black hole.

When the system detects unreachable Packet Forwarding Engine destinations, healing from a traffic black hole is attempted. If the healing fails, the system turns off the interfaces, thereby stopping the black hole and initiating the recovery process.

The recovery process consists of the following steps:

1. Fabric plane restart phase: Healing is attempted by restarting the fabric planes one by one.
2. Fabric plane and FPC restart phase: Healing is attempted by restarting both the fabric planes and the FPCs. If there are bad FPCs that are unable to initiate high-speed links to the fabric after reboot, creation of a traffic black hole is limited because no interfaces are created for these FPCs.
3. FPC offline phase: Traffic black hole is limited by turning the SIBs offline and by turning off interfaces because previous attempts at recovery have failed.

[ [action-fpc-restart-disable](#), [show chassis fabric unreachable-destinations](#), [show chassis fabric reachability](#) ]

- **Forwarding table filter behavior (T Series routers)**—For T Series routers other than T4000, a packet forwarded by the forwarding table reaches the egress forwarding table filter irrespective of whether the packet is actually forwarded by the forwarding table or not. The packet reaches the egress filter even if the route points to reject or discard next hops.

On a T4000 Type 5 FPC, the packet reaches the egress filter only if it is forwarded by the forwarding table.

[[Applying Filters to Forwarding Tables](#)]

---

## High Availability (HA) and Resiliency

- **Support for VRRPv3 (M Series and MX Series routers)**—Junos OS Release 12.2 supports Virtual Router Redundancy Protocol version 3 (VRRPv3). The support for VRRPv3 is implemented in compliance with RFC 5798, *Virtual Router Redundancy Protocol (VRRP) Version 3 for IPv4 and IPv6*. Additionally, Junos OS Release 12.2 supports VRRP MIB for VRRPv3. The support for VRRP MIB for VRRPv3 is implemented in compliance with RFC 6527, *Definitions of Managed Objects for the Virtual Router Redundancy Protocol Version 3 (VRRPv3)*.

To enable VRRPv3, set the **version-3** statement at the **[edit protocols vrrp]** hierarchy level.



**NOTE:** When enabling VRRPv3, you must ensure that VRRPv3 is enabled on all the VRRP routers in the network. This is because VRRPv3 does not interoperate with previous versions of VRRP.

The output of the **show vrrp** command has been modified to indicate the VRRP version that is enabled on the router.

[\[Junos OS Support for VRRPv3\]](#)

- **Support to reduce the VRRP convergence time for quicker traffic restoration (M Series and MX Series routers)**—Enables faster convergence time for the Virtual Router Redundancy Protocol (VRRP), thereby reducing the traffic restoration time to less than 1 second.

To reduce the traffic restoration time, use the following statements at the **[edit protocols vrrp]** hierarchy level:

- **delegate-processing** statement to configure the distributed periodic packet management process (ppmd) to send VRRP advertisements when the ppmmd process is busy.
- **skew-timer-disable** statement to disable the skew timer, thereby reducing the time required to transition to the master state.
- **global-advertisements-threshold** statement to configure the number of fast advertisements that can be missed by a backup router before it starts transitioning to the master state.

You can use the **show protocols vrrp** configuration mode command to see the VRRP configuration information.



**NOTE:**

- The reduction in convergence time is not applicable when VRRP is configured over integrated routing and bridging (IRB) interfaces, aggregated Ethernet interfaces, and multichassis link aggregation group (MC-LAG) interfaces.
- Compared to other routers, the convergence time and the traffic restoration time is less for MX Series routers with MPCs.
- Reduction in convergence time is applicable for all types of configurations at the physical interface but the convergence time might not be less than 1 second for all the configurations. The convergence time depends on the number of groups that are transitioning from the backup to the master state and the interval at which these groups are transitioning.

[\[Improving the Convergence Time for VRRP, Configuring VRRP to Improve Convergence Time\]](#)

## Interfaces and Chassis

---

- **Support for SONET/SDH OC192/STM64 MICs (MX Series routers)**—The following SONET/SDH interface features are supported on the 1-port SONET/SDH OC192/STM64 MIC (model number MIC-3D-IOC192-XFP):
  - SONET or SDH framing—To enable SONET or SDH framing, include the **framing** statement at the `[edit chassis fpc slot-number pic pic-number port port-number]` hierarchy level.
  - Default framing mode is SONET.
  - Total MIC bandwidth cannot exceed 10 Gbps.
  - The single port is configured as clear channel with the speed of OC192 or STM64. The default port speed is OC192.
  - The MIC supports remote and local loopback. Loopbacks can be configured independently on the port.
  - Automatic protection switching (APS) support is based on K1/K2 bytes in SONET frames.
  - The following header bytes can be configured on the MIC using the CLI:
    - Section user channel bytes: F1
    - Line user channel bytes: K1, K2, S1
    - Path user channel bytes: G1, F2, Z3, Z4, C2, E1
  - Configuration for the defect trigger can be ignored. Such ignored defects do not contribute to the interface being marked as down or up.
  - Clock source can be set as external or internal.
  - Path trace identifier to identify the path of the circuit.
  - Incrementing or fixed STM ID to enable interoperability with older equipment.



**NOTE:** The following features are not supported on the 1-port SONET/SDH OC192/STM64 MICs:

- Aggregate SONET (link bundling)
- Multirate configuration
- Link Capacity Adjustment Scheme (LCAS)
- Virtual concatenation

---

### [*Configuring SONET/SDH Physical Interface Properties*]

- **Support for Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX Series routers)**—Junos OS Release 12.2 supports circuit emulation interfaces on MX Series routers. The Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC

with SFP (model number: MIC-3D-4COC3-1COC12-CE) is rate-selectable and can be configured as 4-port OC3/STM1 or 1-port OC12/STM4.



**NOTE:** Junos OS Release 12.2 supports only the rate-selectable 4-port OC3/STM1 MIC.

The following features are supported on the Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP:

- Per-MIC SONET/SDH framing
- Internal and loop clocking
- Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)
- Structure-aware TDM Circuit Emulation Service over Packet-Switched Network (CESoPSN)
- Pseudowire Emulation Edge to Edge (PWE3) control word for use over an MPLS PSN

*[Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Configuring SAToP on 4-port Channelized OC3/STM1 Circuit Emulation MICs, Configuring CESoPSN on 4-Port Channelized OC3/STM1 Circuit Emulation MICs]*

- **Support for centralized clocking on the Enhanced MX Switch Control Board (SCB) (MX240, MX480, and MX960 routers)**—The Enhanced MX SCB uses the centralized Stratum 3 clock module to provide the following features:
  - Clock monitoring, filtering, and holdover
  - Hitless transition from a distributed to centralized clocking mode
  - Distribution of the selected chassis clock source to downstream network elements through supported line interfaces

You can view the centralized clock module information with the command **show chassis synchronization clock-module**.

*[Examples: Configuring Centralized Clocking on the Enhanced MX Switch Control Board]*

- **Support for 100-Gigabit Ethernet MIC with CXP (MIC3-3D-1X100GE-CXP) for the MPC3E on MX240, MX480, and MX960 routers**—The 100-Gigabit Ethernet MIC with CXP (MIC3-3D-1X100GE-CXP) is a 1-port 100-Gigabit Ethernet MIC with a standards-compliant 100GBASE-SR10 interface. The 100-Gigabit Ethernet MIC with CXP uses 100-Gigabit CXP optical transceiver modules for connectivity. It supports up to ten 10-Gigabit Ethernet interfaces and occupies MIC slot 0 or 1 in the MPC3E. The 100-Gigabit Ethernet MIC with CXP supports the same features as the other MICs supported on the MPC3E.

*[MPC3E on MX Series Routers Overview]*

- **Support for 40-Gigabit Ethernet MIC with QSFP+ (MIC3-3D-2X40GE-QSFPP) for the MPC3E on MX240, MX480, and MX960 routers**—The 40-Gigabit Ethernet MIC with QSFP+ (MIC3-3D-2X40GE-QSFPP) is a 2-port 40-Gigabit Ethernet MIC with a

standards-compliant 40GBASE-SR4 interface. It uses quad small form-factor pluggable (QSFP+) optical transceiver modules for connectivity. It occupies slot 0 or 1 in the MPC3E and supports the same features as the other MICs supported on the MPC3E.

[\[MPC3E on MX Series Routers Overview\]](#)

- **Ethernet OAM functionality on MPC3E**—Enables OAM-related operations such as link fault management and link discovery on MPC3E.

The following OAM features are supported on MPC3E:

- Fault detection using continuity check protocol
- Path discovery using link trace protocol
- Fault verification and isolation using loopback protocol
- Distributed PPMD for improved scaling
- GRES support
- RDI support
- Action profiles

[\[MPC3E MIC Overview\]](#)

- **Firewall, Network Address Translation, and intrusion detection service on MPC3E**—Junos OS Release 12.2 supports firewall services, Network Address Translation, and intrusion detection services on MPC3E.

[\[MPC3E MIC Overview\]](#)

- **Dynamic application awareness support on MPC3E**—Adds support for dynamic application awareness functionality and scaling on MPC3E.

[\[MPC3E MIC Overview\]](#)

- **Port-mirroring support on MPC3E**—Adds support for binding up to two port-mirroring instances to the same Packet Forwarding Engine on MPC3E. This enables you to choose multiple mirror destinations by specifying different port-mirroring instances in the filters. You must include the **port-mirror-instance *instance-name*** statement at the **[edit firewall filter *filter-name* term *term-name* then]** hierarchy level. You must also include the **port-mirror-instance *instance-name*** statement at the **[edit chassis fpc *number*]** hierarchy level to specify the FPC to be used.

[\[MPC3E MIC Overview\]](#)

- **Support for 2-port 10-Gigabit Ethernet MIC with XFP on MPC3E (MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.2, MPC3E (MX-MPC3E-3D) supports the 2-port 10-Gigabit Ethernet MIC with XFP (MIC-3D-2XGE-XFP). All features supported by the 2-port 10-Gigabit Ethernet MIC with XFP continue to be supported on the MPC3E.

[\[MPC3E MIC Overview, MPC3E on MX Series Routers Overview\]](#)

- **Support for active flow monitoring features on the MPC3E (MX Series routers with MPC/MIC interfaces)**—The MPC3E supports active flow monitoring features from Junos OS Release 10.4. Flow monitoring versions 5, 8, and 9 support active flow



monitoring. The active flow monitoring features supported are sampling, sampling with templates, sampling per sampling instance, port mirroring, multiple port mirroring, discard accounting, and flow-tap processing.

*[Protocols and Applications Supported by the MX240, MX480, MX960 MPC3E]*

- **Extends support for flow monitoring services to T4000 routers**—Starting with Junos OS Release 12.2R1, the Multiservices 400 PIC on Enhanced Scaling FPC2 supports passive flow monitoring, flow collection, and dynamic flow capture.

*[Passive Flow Monitoring, Flow Collection, Dynamic Flow Capture]*

- **Support for 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ on Type 5 FPC (T4000 routers)**—Starting with Junos OS Release 12.2, the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (model number PF-24XGE-SFPP) is supported on T4000 routers.

The following major software features are supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- Two-to-one oversubscription of traffic in oversubscribed mode.
- Twenty-four 10-Gigabit Ethernet interfaces in oversubscribed mode or 12 ports in line-rate mode.
- All Junos OS configuration commands supported on the existing 10-Gigabit Ethernet LAN/WAN PIC with SFP+.
- The output of the **show interfaces extensive** operational mode command displays preclassification queue counters.
- Line-rate mode operation of first 12 ports can be achieved by the **set chassis fpc fpc-number pic pic-number linerate-mode** command.
- LAN PHY mode and WAN PHY mode support on a per-port basis.
- Aggregated Ethernet is supported only in line-rate mode.
- 4000 logical interfaces per physical interface and 32,000 logical interfaces per chassis.



**NOTE:** Graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) are supported on T4000 routers.

Note that the preclassification is restricted to two traffic classes, and it is not user-configurable. Traffic is classified as control or best effort with non class-of-service (CoS)-aware tail drops of best effort traffic in oversubscribed mode.

For detailed feature support and exceptions, see [24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview](#).

*[show interfaces extensive]*

- **Support for WAN PHY mode on 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (T4000 routers)**—Starting with Junos OS Release 12.2, WAN PHY mode is

supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP), which is plugged into the Type 5 FPC of T4000 routers.

The following WAN PHY features are supported on the 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+:

- WAN PHY mode on a per-port basis
- Insertion and detection of path trace messages
- Ethernet WAN Interface Sublayer (WIS) object

To configure WAN PHY mode on a per-port basis, set the **wan-phy** option for the **framing** statement at the **[edit interface *interface-name*]** hierarchy level.



**NOTE:** When PHY mode changes, interface traffic is disrupted because of port reinitialization.

When WAN PHY mode is configured on an interface, the following SONET options are supported:

- Loopback (local and remote)
- Path trace
- Trigger options

[\[10-Gigabit Ethernet LAN/WAN PIC Overview 24-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview\]](#)

- **Extends support for unicast RPF loose mode (MX Series routers)**—Junos OS Release 12.2 extends support for unicast reverse path forwarding (unicast RPF) loose mode with the ability to discard packets with the source address pointing to the discard interface to MX Series routers. This feature, in conjunction with Remote Triggered Black Hole (RTBH) filtering, provides a mechanism to discard packets from untrusted sources. BGP policies in edge routers ensure that packets with untrusted source addresses have their next hop set to a discard route. When a packet arrives at the router with an untrusted source address, unicast RPF performs a route lookup of the source address. Because the source address route points to a discard next hop, the packet is dropped. This feature is supported on both IPv4 (**inet**) and IPv6 (**inet6**) address families.

To configure unicast RPF loose mode, include the **mode** option in the **rpf-check** statement at the **[edit interfaces]** hierarchy level.

To configure unicast RPF loose mode with the ability to discard packets, you can use the **rpf-loose-mode-discard inet** statement at the **[edit forwarding options]** hierarchy level. Use the **show interfaces extensive** operational mode command to view the packet drops.

[\[Configuring Unicast RPF\]](#)

- **Switch fabric fault management for T4000 routers**—The T4000 router consists of a Switch Interface Board (SIB) with fabric bandwidth double the capacity of the T1600 router. The fabric fault management functionality is similar to that in T1600 routers.

The fabric fault management functionality involves monitoring all high-speed links connected to the fabric and the ones within the fabric core for link failures and link errors. Action is taken based on the fault, and its location. The actions include:

- Reporting link errors in system log files and sending this information to the Routing Engine.
- Reporting link failures at the Flexible Port Concentrator (FPC) or at the SIB and sending this information to the Routing Engine.
- Marking a SIB in **Check** state.
- Moving a SIB into **Fault** state.

The following are the high-level indications of fabric faults that are monitored by Junos OS:

- An SNMP trap is generated whenever a SIB is reported as **Check** or **Fault**.
- **show chassis alarms**—Indicates that a SIB is in **Check** or **Fault** state.
- **show chassis sibs**—Indicates that a SIB is in **Check** or **Fault** state or that a SIB is in **Offline** state when the SIB initializes (this occurs when the SIB does not power on fully).
- **show chassis fabric fpcs**—Indicates whether any fabric links are in error on the FPC's side.
- **show chassis fabric sibs**—Indicates whether any fabric links are in error on the SIB's side.
- The `/var/log/messages` system log messages file at the Routing Engine has error messages with the prefix **CHASSISD\_FM\_ERROR**.
- The SIBs display the **FAIL** LED.

[[Fabric Fault Handling Overview](#), [System Basics: Chassis-Level Features Configuration Guide](#)]

- **Pseudowire TDM MIB support extended to support CESoPSN and SAToP encapsulations on Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX80 routers with a modular chassis, and MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.2, the Pseudowire TDM MIB supports Circuit Emulation Service over Packet-Switched Network (CESoPSN) and Structure-Agnostic TDM over Packet (SAToP) encapsulations configured on Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE). The SAToP encapsulation is supported on T1 and E1 interfaces configured on this MIC. The CESoPSN encapsulation is supported on NxDS0 interfaces configured on the MIC.

[[Interpreting the Enterprise-Specific Pseudowire TDM MIB](#)]

- **IPv6 support for application identification (APPID)**—As of Junos OS Release 12.2, APPID is enabled for IPv6 packets. There is no additional configuration needed for IPv6 support; when AI is enabled, IPv6 support is enabled automatically.
- **Support IPv6 in packet-triggered subscribers and policy control (PTSP), application-aware access list (AACL), and local policy decision function (LPDF)**

**services**—As of Junos OS Release 12.2, PTSP, AACL, and LPDF services, including CLI configuration commands, show commands, and data path processing support IPv6 addressing. All statements that previously accepted only IPv4 addresses, address ranges, and address prefix lengths now also accept IPv6 addresses, address ranges, and address prefix lengths.

For LPDF, you can include the **ipv6-address** and **ipv6-prefix-length** fields at the **[edit services local-policy-decision-function statistics aacl-statistics-profile *profile-name* aacl-fields]** hierarchy level to display them with bulk statistics.



**NOTE:** The IPv6 fields (**ipv6-address** and **ipv6-prefix-length**) are only supported for record-type interim at the **[edit services local-policy-decision-function statistics aacl-statistics-profile *profile-name*]** hierarchy levels; therefore, the fields appear only on the S- (Login) record.

The following operational commands support IPv6 capabilities:

- **show services local-policy-decision-function flows interface *interface-name***—The IPv4 flows are shown as currently defined. A section is added that shows the IPv6 flows.
- **show services application-aware-access-list flows**—The IPv4 flows are shown as currently defined. A section is added that shows the IPv6 flows.
- **show services local-policy-decision-function statistics**—The IPv4 flows are shown as currently defined. A section is added that shows the IPv6 flows.
- **show services application-aware-access-list statistics**—The same output is shown for **show services local-policy-decision-function statistics**.
- **show services subscriber sessions**—The IP address entered in this command can be either IPv4 or IPv6.
- **show services subscriber flows**—The command displays the IPv4 or IPv6 address.
- **show services subscriber bandwidth**—The IP address entered in this command can be either IPv4 or IPv6.

[[Configuring AACL Rules](#), [Configuring Statistics Profiles](#)]

- **Support for link fault management (IEEE 802.3ah) features and connectivity fault management (IEEE 802.1ag) protocols on T4000 routers**—Starting with Junos OS Release 12.2, the link fault management features and the connectivity fault management protocols listed in [Table 2 on page 117](#) are supported on T4000 routers with the following PICs:
  - 100-Gigabit Ethernet PIC with CFP (PD-10GE-CFP)
  - 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-10XGE-SFPP)
  - 24-port 10-Gigabit Ethernet PIC (PF-24XGE-SFPP)
  - 10-Gigabit Ethernet LAN/WAN PIC with XFP (PD-4XGE-XFP)
  - 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PD-5-10XGE-SFPP)

Table 2: Link Fault Management Features and Connectivity Fault Management Protocols

Link Fault Management Features	Connectivity Fault Management Protocols
<ul style="list-style-type: none"> <li>• Link discovery</li> <li>• Fault detection</li> <li>• Action profiles</li> <li>• Event thresholds</li> </ul>	<ul style="list-style-type: none"> <li>• Continuity check protocol</li> <li>• Loopback protocol</li> <li>• Linktrace protocol</li> </ul>



**NOTE:** The remote loopback feature mentioned in section 57.2.11 of IEEE 802.3ah is not supported on T4000 routers.

[[IEEE 802.3ah OAM Link-Fault Management Overview](#), [IEEE 802.1ag OAM Connectivity Fault Management Overview](#)]

- **Extends support for port mirroring with next-hop groups to T4000 Type 5 FPC (T4000-FPC5-3D)**—Junos OS Release 12.2 supports port mirroring and multipacket port mirroring on the T4000 Type 5 FPC.

[[Port Mirroring](#)]

- **Support for reporting total statistics for IPv6 traffic traversing through T4000 routers**—Starting with Junos OS Release 12.2, total statistics (sum of local and transit traffic) is reported for traffic traversing through the following PICs on T4000 routers:
  - 10-Gigabit Ethernet IQ2 PIC with XFP (PC-1XGE-TYPE3-XFP-IQ2)
  - 10-Gigabit Ethernet Enhanced IQ2 (IQ2E) PIC with XFP (PC-1XGE-TYPE3-XFP-IQ2E)
  - 12-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-12XGE-SFPP)
  - 24-port 10-Gigabit Ethernet LAN/WAN PIC with SFP+ (PF-24XGE-SFPP)

You can view the IPv6 statistics by issuing the following commands:

- **show snmp mib walk ipv6IfStatsTable**
- [[show interfaces extensive](#)]
- **Merge and sort of VLAN ranges**—Starting with Release 12.2, Junos OS provides a merge and sort feature for VLAN ranges. The merge feature enables Junos OS to merge overlapping VLAN ranges and to display combined values for such ranges. For example, if your configuration has multiple VLAN ranges of 1–15, 12–22, and 17–30, Junos OS displays the VLAN member range as 1–30, which is the combined value for all three ranges. The sort feature enables Junos OS to sort the VLAN range values in such a way that the numeric values are listed in ascending order followed by alphanumeric values. However, if there are only alphanumeric values in the configuration, such values are displayed in the same sequence as they were configured.

Junos OS Release 12.2 and later support the merge and sort feature for the following configuration statements :

- **vlan-id-list** at the [**interfaces *interface-name* unit *unit***] and [**bridge-domains *domain***] hierarchy levels.

- **vlan tags inner-list** at the **[interfaces *interface-name* unit *unit*]** hierarchy level.

[\[Merge and Sort Support for VLAN Ranges\]](#)

- **Single-core Routing Engine support for M7i and M10i routers**—Starting with Junos OS Release 12.2, a single-core Routing Engine is added to the M7i and M10i routers. This Routing Engine is based on the single-core Intel Xeon CPU, operating at 1.73 GHz with 2 MB cache and has two DDR3 DIMM slots operating at 800 MHz that support 4 GB memory with error checking and correction (ECC). The new Routing Engine also supports:
  - 82,574 Gigabit Ethernet Controller
  - 4 GB CompactFlash card
  - USB 2.0
  - Front accessible 64 GB solid-state drive (SSD)

All CLI commands supported on the older Routing Engine are supported on the new Routing Engine.

[\[Supported Routing Engines by Chassis\]](#)

- **Support for hierarchical schedulers on aggregated Ethernet interfaces (MX Series routers)**—Enables you to apply hierarchical schedulers on aggregated Ethernet bundles through the use of interface sets. This feature is supported at egress only on MX Series routers.

You can configure interface sets for aggregated Ethernet interfaces created under static configuration, as well as dynamic configurations.

You can configure class-of-service parameters on aggregated interfaces, in either link-protect or non-link-protect mode. You can configure these parameters at the physical, interface set, and logical interface levels. The CoS configuration is fully replicated for all aggregated Ethernet member links.

You can control the way these parameters are applied by configuring the aggregated interface to operate in scaled mode or replicate mode.

The link membership list and scheduler mode of the interface set is inherited from the underlying aggregated Ethernet interface over the interface set is configured. When an aggregated Ethernet interface operates in link protection mode, or if the scheduler mode is set to **member-link-scheduler replicate**, the scheduling parameters of the interface set are copied to each of the member links.

If the scheduler mode of the aggregated Ethernet interface is set to **member-link-scheduler scale**, the scheduling parameters are scaled based on the number of active member links (scaling factor is  $1/A$ , where  $A$  is the number of active links in the bundle) and applied to each of the aggregated interface member links.

To configure an interface set, include the **interface-set *interface-set-name*** statement at the **[edit class-of-service interfaces]** hierarchy level.

To apply scheduling and queuing parameters to the interface set, include the **output-traffic-control-profile *profile-name*** statement at the **[edit class-of-service interfaces interface-set *interface-set-name*]** hierarchy level.

To apply the traffic control profile to the interface or interface set, include the **output-traffic-control-profile-remaining *profile-name*** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level or the **[edit class-of-service interfaces *interface-name* interface-set *interface-set-name*]** hierarchy level.

[\[Hierarchical Schedulers on Aggregated Ethernet Interfaces Overview\]](#)

- **Support for trunk port enhancements (MX Series routers with MPC3E)**—Extends support of the trunk port features to MX240, MX480, and MX960 routers with MPC3E (model no: MX-MPC3E-3D). You can configure a single logical interface to support a list of VLANs or to accept packets with no VLAN tag. You can also configure multiple logical trunk interfaces on a single physical interface.

You can also configure dynamic profiles for VPLS pseudowires, VLAN identifier translation, and automatic bridge domain configuration. To configure dynamic profiles, include the **profile-name** statement at the **[edit dynamic-profiles]** hierarchy level.

With the VLAN translation feature, you can configure a trunk port interface to translate the VLAN identifier associated with the ingress interface into the VLAN identifier of the destination bridge domain at egress. To configure multiple bridge domains, include the **vlan-id-list** and **vlan-id-range** statements at the **[edit bridge-domains bridge-domain-name]** hierarchy level.

Layer 3 multicast is supported on Layer 2 trunk ports through integrated routing and bridging (IRB) interfaces.

[\[Dynamic Profiles for VPLS Pseudowires\]](#)

[\[Example: Configuring Multiple Bridge Domains with a VLAN ID List\]](#)

- **Support for redundant fabric made on active control boards of MX Series routers**—An FPC working with reduced fabric bandwidth can affect the rerouting process and can cause partial traffic black holes. You can enable increased fabric bandwidth of active control boards for optimal and efficient performance and traffic handling. On an MX960, MX480, or MX240 router, you can configure the active control board to be in redundancy mode or in increased fabric bandwidth mode. In increased fabric bandwidth mode, the maximum number of available fabric planes are used for MX Series routers with Trio chips and the MPC3E. On MX960 routers with active control boards, six active planes are used, and on MX240 and MX480 routers with active control boards, eight active planes are used.

To configure redundancy mode for the active control board, use the **redundancy-mode redundant** statement at the **[edit chassis fabric]** hierarchy level. When you configure this option, all the FPCs use four fabric planes as active planes, regardless of the type of the FPC. If you do not configure this option, increased fabric bandwidth mode is enabled by default on MX Series routers with Switch Control Board (SCB). The MX Series routers that contain the enhanced Switch Control Board (SCB) with Trio chips and the MPC3E, the control boards operate in redundancy fabric mode (all the FPCs use four fabric planes as active planes) by default.

To configure increased bandwidth mode for the active control board, use the **increased-bandwidth** statement at the **[edit chassis fabric]** hierarchy level. When you configure this option, 6 active planes are used.

Configuring this feature does not affect the system. You can configure this feature without restarting the FPC or restarting the system.

You can use the **show chassis fabric redundancy-mode** command to verify whether the redundancy fabric mode is enabled.

[[System Basics: Chassis-Level Features Configuration Guide](#)]

- **Interoperability of Type 3 FPCs and Type 4 FPCs with Type 5 FPCs (T4000 routers)**—Support for interoperability of T640 Enhanced Scaling FPC3, T1600 Enhanced Scaling FPC4, and T640 Enhanced Scaling FPC4-1P with T4000 FPC5 is possible with fabric notification translation. This feature is supported on T4000 routers.

Basic packet forwarding, IPv4, IPv6, MPLS, and multicast (data plane) are currently supported through this feature.

[[T4000 Core Router PIC Guide](#)]

- **Support for user-defined system identifier in LACP**—The user-defined system identifier in the Link Aggregation Control Protocol (LACP) enables two ports from two separate routers (M Series or MX Series routers) to act as though they were part of the same aggregate group.

[[Configuring Aggregated Ethernet LACP](#)]

- **SAToP support extended to MIC-3D-4COC3-1COC12-CE**—Starting with Junos OS Release 12.2R1, the support for Structure-Agnostic time-division multiplexing over Packet (SAToP) is extended to MIC-3D-4COC3-1COC12-CE. You can configure 84 T1 channels on each coc3 interface on this MIC.

[[Configuring SAToP on 4-port Channelized OC3/STM1 Circuit Emulation MICs](#)]

- **SONET/SDH support on the Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE)**—Starting with Junos OS Release 12.2R1, the SONET/SDH interfaces are supported on the Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP.

You can configure the following SONET/SDH physical interface properties:

- Loopback capability
  - Framing
  - Interface speed
  - Automatic Protection Switching
  - External or loop timing
  - Internal timing
  - Up or down defect hold-time
- **Support for PWE3 routing extension in CESoPSN for LDP/RSVP signaling**—Support for PWE3 routing extension in CESoPSN for LDP/RSVP signaling is available in Junos OS Release 12.2R1 and later releases.

[[Configuring the Pseudowire Interface](#)]



- **CESoPSN support for MIC-3D-4COC3-1COC12-CE, for all values of *N* in *NxDSO* interfaces (MX80, MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.2, Circuit Emulation Service over Packet-Switched Network (CESoPSN) is supported on the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC (MIC-3D-4COC3-1COC12-CE) for all values of *N* in *NxDSO* interfaces. CESoPSN encapsulation is supported on *NxDSO* interfaces.



**NOTE:** The Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC (MIC-3D-4COC3-1COC12-CE) supports CESoPSN services without channel-associated signaling (CAS).

An *NxDSO* interface can be configured from either a channelized T1 interface (CT1) or a channelized E1 interface (CE1).

The value of *N* is 24 when a DSO interface is configured from a CT1 interface and 31 when a DSO interface is configured from a CE1 interface.

To configure an *NxDSO* interface, configure the **set ct1-x/y/z:1:3 partition 1 timeslot 1-4, 9, 22-24 interface-type ds** statement at the [edit interfaces] hierarchy level. Then create the DS interface by configuring the **set ds-x/y/z:1:3:1 encapsulation cesopsn unit 0** statement at the [edit interfaces] hierarchy level.

[\[Configuring CESoPSN on 4-Port Channelized OC3/STM1 Circuit Emulation MICs\]](#)

- **802.1ad provider bridge support**—Extends support for 802.1ad provider bridge features to MX Series MPC3E interfaces.

[\[Configuring and Applying IEEE 802.1ad Classifiers\]](#)

- **Support to interoperate the 100-Gigabit Ethernet PIC on Type 4 FPC (T1600 routers) with the 100-Gigabit Ethernet PIC on Type 5 FPC (T4000 routers)**—Enables the interoperability between the 100-Gigabit Ethernet PIC on Type 4 FPC (on T1600 routers) and the 100-Gigabit Ethernet PIC on Type 5 FPC (on T4000 routers) by enabling a source address (SA) multicast bit steering mode on the 100-Gigabit Ethernet PIC on Type 5 FPC. The SA multicast bit steering mode uses the multicast bit in the source MAC address for packet steering.

By default, the SA multicast bit steering mode is not enabled on the 100-Gigabit Ethernet PIC on Type 5 FPC. To enable the SA multicast bit steering mode on the 100-Gigabit Ethernet PIC on Type 5 FPC, include the **forwarding-mode sa-multicast** statement at the [edit chassis fpc fpc-slot-number pic pic-slot-number] hierarchy level.



**NOTE:** The configuration of the forwarding-mode sa-multicast statement results in a PIC bounce—that is, the 100-Gigabit Ethernet PIC on Type 5 FPC goes offline and comes back online.

[\[Interoperability Between the 100-Gigabit Ethernet PIC on Type 4 FPC and the 100-Gigabit Ethernet PIC on Type 5 FPC, Configuring the Interoperability Between the 100-Gigabit Ethernet PIC on Type 5 FPC and the 100-Gigabit Ethernet PIC on Type 4 FPC\]](#)

- **Support for Precision Time Protocol (MX80, MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.2, Precision Time Protocol (PTP), also known as IEEE 1588v2, is supported on MX80 routers with precision timing support (MX80-P). On MX240, MX480, and MX960 routers, PTP is supported on the Enhanced Module Port Concentrator (MPCE) model number MX-MPC2E-3D-P and its Ethernet Modular Interface Cards (MICs).

PTP synchronizes clocks between nodes in a network, thereby enabling the distribution of an accurate clock over a packet-switched network. This synchronization is achieved through packets that are transmitted and received in a session between a master clock and a slave clock.



**NOTE:** Unified in-service software upgrade (unified ISSU) is currently not supported when PTP is configured on MX240, MX480, and MX960 routers.



**NOTE:** To switch between the PTP and Synchronous Ethernet modes, you must first deactivate the configuration for the current mode and then commit the configuration. Wait for a short period of 30 seconds, configure the new mode and its related parameters, and then commit the configuration.

[[System Basics: Chassis-Level Features Configuration Guide, PTP Operational Mode Commands](#)]

- **Support for Synchronous Ethernet and Precision Time Protocol on MX Series routers with Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MX Series routers)**—Starting with Junos OS Release 12.2R1, Synchronous Ethernet and Precision Time Protocol (PTP) are supported on MX Series routers with Channelized OC3/STM1 (Multi-Rate) Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE). The clock derived by Synchronous Ethernet and PTP is used to drive the SONET/SDH interfaces on this MIC.

[[System Basics: Chassis-Level Features Configuration Guide](#)]

- **Support for the combined operation of Synchronous Ethernet and Precision Time Protocol or hybrid mode (MX Series 3D Universal Edge Routers)**—Combined operation of Synchronous Ethernet and Precision Time Protocol (PTP), also known as hybrid mode, is supported on the MX80 routers with precision timing (MX80-P) and with timing (MX80-T). It is also supported on MX240, MX480, and MX960 routers. On the MX240, MX480, and MX960 routers, the combined operation is possible only when the PTP client and the Synchronous Ethernet source are on the same enhanced MPC and are traceable to the same master clock. In hybrid mode, the Synchronous Ethernet equipment clock (EEC) on the Modular Port Concentrator (MPC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP (also known as IEEE 1588v2) for time synchronization.



**NOTE:** When acting as PTP slaves, MX80-P routers can accept any external Synchronous Ethernet clock as reference and do not support building-integrated timing supply (BITS) input as frequency source in hybrid mode. Only Synchronous Ethernet sources are allowed in hybrid mode.

Synchronous Ethernet and PTP provide frequency and phase synchronization; however, the accuracy in the order of nanoseconds is difficult to achieve through PTP or Synchronous Ethernet, and they do not support a large number of network hops. Hybrid mode resolves these issues by extending the number of network hops and also provides the clock synchronization accuracy in the order of tens of nanoseconds.

To configure hybrid mode, include the **hybrid synchronous-ethernet-mapping clock-source *ip-address* interface *interface-name1*** statement at the **[edit protocols ptp slave]** hierarchy level.

To set the Ethernet Synchronization Message Channel (ESMC) from the PTP clock class, include the **convert-clock-class-to-quality-level** statement at the **[edit protocols ptp slave]** hierarchy level.

To override the default PTP clock class to ESMC mapping, include the **clock-class-to-quality-level-mapping quality-level *ql-value* clock-class *clock-class-value*** statement at the **[edit protocols ptp slave]** hierarchy level, where **clock-class** indicates the current state of the clock and the **quality-level** indicates the clock type.

Note that when the selected Synchronous Ethernet reference fails, the router continues to work in PTP mode. You can use the **show ptp hybrid status** operational command to find the current operating mode.

Unified in-service software upgrade (unified ISSU) is not supported when clock synchronization is configured for hybrid mode on MX80-P routers and MX80-T routers, and on the MICs and enhanced MPCs on MX240, MX480, and MX960 routers.



**NOTE:** To switch between the PTP and Synchronous Ethernet modes, you must first deactivate the configuration for the current mode and then commit the configuration. Wait for 30 seconds, configure the new mode and its related parameters, and then commit the configuration.

- **MAC address validation in enhanced network services modes (MX Series routers)**—MAC address validation is optimized for scaling when the router is configured for Enhanced IP Network Services mode or Enhanced Ethernet Network Services mode. When MAC address validation is enabled, the router compares the IP source and MAC source addresses against trusted addresses, and forwards or drops the packets according to the match and the validation mode. This feature is not available for IPv6.



**NOTE:** When the router is configured for either of the enhanced network services modes, MAC address validation is supported only on MPCs. If the router has both DPCs and MPCs, or only DPCs, you cannot configure the chassis to be in enhanced mode.

In contrast, when the router is configured for a normal (non-enhanced) network services mode, MAC address validation is supported on both DPCs and MPCs. The router can be populated completely with one or the other type of line card, or have a mix of both types. Normal network services mode is the default.

To configure an enhanced network services mode, include the **network-services service** statement at the **[edit chassis]** hierarchy level, and then configure MAC address validation as usual.



**NOTE:** In normal network services mode, you can use the **show interfaces statistics interface-name** command to display a per-interface count of the packets that failed validation and were dropped. In enhanced network services modes, this command does not count the dropped packets; you must contact Juniper Networks Customer Support for assistance in collecting this data.

- **Fail filters for RPF checks in dynamic profiles (MX Series routers)**—By default, unicast RPF checks prevent DHCP packets from being accepted on interfaces protected by the RPF check. When you enable an RPF check with a dynamic profile, you must configure a fail filter that identifies and passes DHCP packets.

To configure a fail filter, include the **fail-filter filter-name** statement at the **[edit dynamic-profiles profile-name interfaces interface-name unit logical-unit-number family family rpf-check]** hierarchy level. To configure the terms of the fail filter, include the **filter filter-name** statement at the **[edit firewall family family]** hierarchy level. Include conditions in a filter term to identify DHCP packets, such as **from destination-port dhcp** and **from destination-address 255.255.255.255/32**. Define another filter term to drop all other packets that fail the RPF check. This feature is available for both IPv4 and IPv6 address families.

To confirm that the fail filter is active, issue the **show subscribers extensive** command, which displays the name of active filters.

- **Setting the 802.1p field for host-generated traffic (MX Series routers)**—On MPCs and Enhanced Queuing DPCs, you can configure the IEEE 802.1p bits in the 802.1p field—also known as the Priority Code Point (PCP) field—in the Ethernet frame header for host outbound packets (control plane traffic). In releases earlier than 12.2R3, this field is not configurable; instead it is set by CoS automatically for host outbound traffic.

To configure a global default value for this field for all host outbound traffic, include the **default value** statement at the **[edit class-of-service host-outbound-traffic]**

**ieee-802.1p** hierarchy level. This configuration has no effect on data plane traffic; you configure rewrite rules for these packets as always.

You cannot configure a default value for the 802.1p bits for host outbound traffic on a per-interface level. However, you can specify that the CoS 802.1p rewrite rules already configured on egress logical interfaces are applied to all host outbound packets on that interface. To do so, include the **rewrite-rules** statement at the **[edit class-of-service host-outbound-traffic ieee-802.1p]** hierarchy level. This capability enables you to set only the outer tags or both the outer and the inner tags on dual-tagged VLAN packets. (On Enhanced Queuing DPCs, both inner and outer tags must be set.)

This feature includes the following support:

- Address families—IPv4 and IPv6
- Interfaces—IP over VLAN demux, PPP over VLAN demux, and VLAN over Gigabit Ethernet
- Packet types—ARP, ANCP, DHCP, ICMP, IGMP, and PPP
- VLANs—Single and dual-tagged
- **Improvements to interface transmit statistics reporting (MX Series routers)**—On MX Series routers, the logical interface-level statistics show only the offered load, which is often different from the actual transmitted load. To address this limitation, Junos OS introduces a new configuration statement in Releases 11.4R3, 12.1R4, and 12.2R3 and later. The new configuration statement, **interface-transmit-statistics** at the **[edit interface interface-name]** hierarchy level, enables you to configure Junos OS to accurately capture and report the transmitted load on interfaces.

When the **interface-transmit-statistics** statement is included at the **[edit interface interface-name]** hierarchy level, the following operational mode commands report the actual transmitted load:

- **show interface interface-name <detail | extensive>**
- **monitor interface interface-name**
- **show snmp mib get objectID.ifIndex**



**NOTE:** This configuration is not supported on Enhanced IQ (IQE) and Enhanced IQ2 (IQ2E) PICs.

The **show interface interface-name** command also shows whether the **interface-transmit-statistics** configuration is enabled or disabled on the interface.

[Class of Service]

- **Optical transceiver support for MIC3-3D-2X40GE-QSFPP on MPC3E (MX240, MX480, and MX960 routers)**—Starting with Junos OS Release 12.2, the 2-port 40-Gigabit Ethernet MIC with QSFPP (MIC3-3D-2X40GE-QSFPP) on MPC3E supports the QSFPP-40GBase-LR4 optical transceiver.

*[MX Series 3D Universal Edge Routers Line Card Guide]*

- **PICs supported on T4000 Core Routers**—Starting with Junos OS Release 12.2 R2, the following PICs are supported on T4000 routers.
  - Enhanced Scaling FPC2 (model number: T640-FPC2-ES) on the T4000 router supports the following Type 2 PICs:
    - ATM2 OC12/STM4 IQ PIC (model number: PB-2OC12-ATM2-SMIR)
    - Channelized OC12/STM4 Enhanced IQ (IQE) PIC with SFP (model number: PB-4CHOC12-STM4-IQE-SFP)
    - Channelized OC48/STM16 Enhanced IQ (IQE) PIC with SFP (model number: PB-1CHOC48-STM16-IQE-SFP)
    - Gigabit Ethernet PICs with SFP (model numbers: PB-2GE-SFP and PB-4GE-SFP)
    - Gigabit Ethernet IQ PIC with SFP (model number: PB-2GE-SFP-QPP)
    - Gigabit Ethernet IQ2 PIC with SFP (model number: PB-8GE-TYPE2-SFP-IQ2)
    - SONET/SDH OC12/STM4 (Multi-Rate) PIC with SFP (model number: PB-4OC3-4OC12-SON-SFP)
    - SONET/SDH OC48c/STM16 EOL PIC (model number: PB-1OC48-SON-SMSR)
    - SONET/SDH OC12c/STM4 EOL PICs (model numbers: PB-4OC12-SON-MM and PB-4OC12-SON-SMIR)
    - SONET/SDH OC48c/STM16 EOL PIC with SFP (model number: PB-1OC48-SON-SFP)
    - SONET/SDH OC48/STM16 (Multi-Rate) PIC with SFP (model number: PB-1OC48-SON-B-SFP)
    - SONET/SDH OC3/STM1 (Multi-Rate) PIC with SFP (model number: PB-4OC3-1OC12-SON2-SFP)
    - Tunnel Services PICs (model number: PB-TUNNEL)
  - Enhanced Scaling FPC3 (model number: T640-FPC3-ES) on the T4000 router supports the following Type 3 PIC:
    - 10-Gigabit Ethernet PIC with XENPAK (model number: PC-1XGE-XENPAK)

[See [T4000 PICs Supported.](#)]

## Junos OS Installation and Upgrade

- **Licensable ports on MX5, MX10, and MX40 routers**—License keys are available to enhance the port capacity on MX5, MX10, and MX40 routers up to the port capacity of an MX80 router. The MX5, MX10, and MX40 routers are derived from the modular MX80 chassis with similar slot and port assignments, and provide all functionality available on an MX80 router, but at a lower capacity. Restricting port capacity is achieved by making a set of MIC slots and ports licensable. MICs without a license are locked, and are unlocked or made usable by installing appropriate upgrade licenses.

[[Junos OS License Key](#)]

- **Release-based capacity licenses in chassis mode**—Support for enforcing license-based restrictions while upgrading Junos OS is provided, along with support for an upgrade license key for license-based features. When upgrading a Junos OS installation, a license for a feature is considered valid if the release version in the license key is greater than or equal to the release version of the software upgrade. Valid license keys are displayed in the **show system license** command output.

[[Junos OS License Key](#)]

- **Smooth upgrade and downgrade procedures for T4000 routers**—You can upgrade a T1600 router with SF-based SIB (SIB-I8-SF with model number—SIB-I-T1600-S) and a T640 router with F16 2.0-based SIB only (SIB-I8-F16 2.0 SIBs with model number—SIB-I-T640-B-S) to a T4000 router. You can also downgrade a T4000 router to a T640 router or a T1600 router.

To upgrade from a T640 chassis or T1600 chassis to a T4000 chassis, use the **set chassis fabric upgrade-mode t4000** command. To downgrade from a T4000 chassis to a T1600 chassis or a T640 chassis, use the **set chassis fabric upgrade-mode default** command.

[[T4000 Core Router Hardware Guide](#)]

- **ICMP redirect**—As of Junos OS Release 12.2, ICMP redirect messages for the IPv6 family are enabled by default for devices using line cards with the Junos Trio chipset. This feature checks all IPv6 packets that enter and exit on the interface and provides ICMP redirect messages to notify hosts when a better route is available for a particular destination. All redirects can be disabled by using the **set system no-redirects** command.
- **Retain, delete, or validate add-on packages during installation**—Three operational statements allow you to retain, delete, or validate a set of software add-on packages when upgrading or downgrading a Junos OS software package. This allows you to manage multiple software add-on packages at the same time. The commands are **request system software add set**, **request system software delete set**, and **request system software validate set**.

[[Upgrading Software Packages](#)]

## Junos OS XML API and Scripting

---

- **libslax distribution supports SLAX script development**—libslax is an open-source implementation of the SLAX language using the "New BSD License." libslax is written in C and is built on top of the libxml2, libxslt, and libexslt libraries. The libslax distribution contains the libslax library, incorporates a SLAX writer and SLAX parser, a debugger, a profiler, and the SLAX processor (slaxproc). The SLAX processor is a command-line tool that can validate SLAX script syntax, convert between SLAX and XSLT formats, and format, debug, or run SLAX scripts.

The libslax tools are included as part of the standard Junos OS. However, you can download and install the libslax distribution on a computer with a UNIX-like operating system to develop SLAX scripts outside of Junos OS. Links to the current releases, source code, documentation, and support materials for libslax are available at the SLAX community and support site at <http://www.libslax.org>.

[[libslax Distribution Overview](#)]

- **Support for NETCONF tracing operations**—Starting with Junos OS Release 12.2, you can configure tracing operations for the NETCONF XML management protocol. NETCONF tracing operations record NETCONF session data in a trace file. The default trace file is `/var/log/netconf`. By default, NETCONF tracing operations are not enabled.

You configure NETCONF tracing operations at the **[edit system services netconf traceoptions]** hierarchy level. To enable NETCONF tracing operations and to trace all incoming and outgoing data from NETCONF sessions on that device, configure the **flag all** statement. To restrict tracing to only incoming or outgoing NETCONF data, configure the flag value as either **incoming** or **outgoing**, respectively. Additionally, to restrict the trace output to include only those lines that match a particular expression, configure the **file match** statement and define the regular expression against the matched output.

To control the tracing operation from within a NETCONF session, configure the **on-demand** statement. This requires that you start and stop tracing operations from within the NETCONF session. If you configure the **on-demand** statement, you must issue the `<rpc><request-netconf-trace><start/></request-netconf-trace></rpc>` RPC in the NETCONF session to start tracing operations for that session. To stop tracing for that NETCONF session, issue the `<rpc><request-netconf-trace><stop/></request-netconf-trace></rpc>` RPC.

[[Example: Configuring NETCONF Tracing Operations](#)]

- **jcs:load-configuration template supports the rollback parameter and a null configuration**—The **jcs:load-configuration** template supports the **rollback** parameter, which rolls back the configuration to a previously committed configuration. Specify the rollback number of the configuration, and the configuration is loaded from the associated file.

The **jcs:load-configuration** template accepts a NULL configuration for the **configuration** parameter. If you supply a NULL configuration, the template performs a simple commit of the candidate configuration. Otherwise, configuration changes are incorporated into the candidate configuration as specified by the **action** parameter.



- **jcs:open() extension function support for routing-instances**—The **jcs:open()** extension function returns a connection handle that is used to execute RPCs on a local or remote device. To redirect the SSH connection to originate from within a specific routing instance, include the name of the routing instance in the connection parameters. The routing instance must be configured at the **[edit routing-instances]** hierarchy level, and the remote device must be reachable either using the routing table for that routing instance or from one of the interfaces configured under that routing instance.

*[open() Function (jcs Namespace)]*

- **Support for commit script access to the pre-inheritance candidate configuration in configure private sessions**—Commit scripts can invoke the **<get-configuration>** RPC in a private configuration session to retrieve the private, pre-inheritance candidate configuration for that session. The **<get-configuration>** RPC includes the **database-path** attribute, which is used to specify the location of the pre-inheritance configuration database. In addition, the global variable, **\$junos-context** contains the **commit-context/database-path** element that stores the location of the session's pre-inheritance candidate configuration.

To construct a commit script that retrieves the pre-inheritance candidate configuration specific to that session, include the **<get-configuration>** RPC in the commit script, and set the **<database-path>** attribute to **\$junos-context/commit-context/database-path**.

## Layer 2 Ethernet Services

- **Support for Layer 2 features on the SONET/SDH OC192/STM64 MICs (MX Series routers)**—The following Layer 2 features are supported on the 1-port SONET/SDH OC192/STM64 MIC (model number MIC-3D-1OC192-XFP):
  - Interface MTU settings (range: 256–9192 bytes).
  - High-Level Data Link Control (HDLC) payload scrambling.
  - HDLC CRC checking supports two modes—**crc-16** and **crc-32**.
  - Default idle cycle transmit value is 0x7E.
  - Encapsulations:
    - **cisco-hdlc**—Cisco-compatible HDLC framing
    - **cisco-hdlc-ccc**—Cisco-compatible HDLC framing for a cross-connect
    - **cisco-hdlc-tcc**—Cisco-compatible HDLC framing for a translational cross-connect
    - **flexible-frame-relay**—Multiple Frame Relay encapsulations
    - **frame-relay**—Frame Relay encapsulation
    - **frame-relay-ccc**—Frame Relay for a cross-connect
    - **frame-relay-tcc**—Frame Relay for a translational cross-connect
    - **ppp**—Serial Point-to-Point Protocol (PPP) device
    - **ppp-ccc**—Serial PPP device for a cross-connect
    - **ppp-tcc**—Serial PPP device for a translational cross-connect

- MPLS circuit cross-connect
- MPLS translational cross-connect
- MPLS fast reroute



**NOTE:** The following Layer 2 encapsulations are not supported on the 1-port SONET/SDH OC192/STM64 MICs:

- Multilink Frame Relay end-to-end (FRF.15)
- Multilink Frame Relay end-to-end (FRF.16)
- Multilink PPP
- Generic framing procedure (GFP)

*[encapsulation (Physical Interface), MTU, Configuring SONET/SDH HDLC Payload Scrambling]*

- **Subscriber Secure Policy support for Layer 2 Tunneling Protocol (L2TP) subscribers (MX Series routers)**—Subscriber Secure Policy supports L2TP subscribers terminating at the Layer 2 network server.

*[Subscriber Secure Policy and L2TP LNS Subscribers]*

- **Extends support for MAC filtering, accounting, policing, and learning to T4000 Type 5 FPC (T4000-FPC5-3D)**—Support for logical interface-level MAC filtering, accounting, policing, and learning for source media access control (MAC) is extended to the T4000 Type 5 FPC. The following features are not supported on the T4000 Type 5 FPC:

- MAC filtering, accounting, and policing for destination MAC at the logical interface level.



**NOTE:** Because destination MAC filtering is not supported, the hardware is configured to accept all multicast packets. This configuration enables the OSPF protocol to work.

- Premium MAC policers at the logical interface level.
- MAC filtering, accounting, and policing at the physical interface level.

*[12-port 10-Gigabit Ethernet LAN/WAN PIC on Type 5 FPC Overview 100-Gigabit Ethernet PIC on Type 5 FPC Overview]*

- **MIB support for Layer 2 policer statistics (MX Series routers)**—Adds MIB functionality to display Layer 2 policer statistics on MX Series routers. Use the **show interface interface-name detail** command to view Layer 2 policer statistics in the MIB.

*[show snmp mib]*

- **Support for Layer 2 and Layer 2.5 features (MX Series routers with MPC3E)**—Starting with Junos OS Release 12.2, support for the Layer 2 and Layer 2.5 protocols is extended

to MX240, MX480, and MX960 routers with MPC3E (model number MX-MPC3E-3D). The following features are supported:

- IGMP snooping for multichassis link aggregation group (MC-LAG) interfaces  
*[[IGMP Snooping in MC-LAG Active-Active on MX Series Router Overview](#)]*
- Configurable label block sizes for VPLS
- Connectivity fault management process flooding to interfaces based on mesh groups
- Layer 2 address learning in logical systems  
*[[Layer 2 Learning and Forwarding in a Logical System Overview](#)]*
- Ethernet Ring Protection Switching for multiple ring instances on the same physical ring  
*[[Ethernet Ring Protection Using Ring Instances for Load Balancing](#)]*
- Transit and bypass static label-switched paths (LSPs)
- Layer 2 Gigabit Ethernet logical interface policing
- Static LSP statistics
- Multiple VLAN Registration Protocol (MVRP)—IEEE 802.1ak-2007  
*[[Understanding Multiple VLAN Registration Protocol \(MVRP\) on MX Series Routers](#)]*

- **Support for Layer 2 Ethernet OAM (MX Series routers with MPC3E)**—Extends support for Layer 2 Ethernet OAM features (802.3ah only) through Junos OS Release 12.2 to MX240, MX480, and MX960 routers with MPC3E (model number MX-MPC3E-3D).

The following Layer 2 Ethernet OAM functions are supported:

- Distributed periodic packet management process (ppmd) for improved scaling
- Graceful Routing Engine switchover (GRES)
- Remote defect indication (RDI)
- Configuration of action profiles

*[[IEEE 802.3ah OAM Link-Fault Management Overview](#)]*

- **Layer 2 protocols on MPC3E**—Enables Layer 2 protocols on MPC3E.

The following Layer 2 protocols are supported on MPC3E:

- BPDU-protect

The BPDU-protect feature is part of an L2CPD module for MX Series devices that runs the spanning tree suite of protocols. Spanning Tree Protocols (STPs) break loops in a Layer 2 bridged network, protecting the network from possible broadcast storms. BPDU-protect helps prevent misbehaving applications or devices from interfering with STP operations.

- Root guard

The root guard feature protects the root bridge by restricting the core bridge from allowing the edge bridge to declare any others as “parent.” This ensures that the core bridge is always elected as a root bridge and protected.

- BPDU loop protect

STP breaks loops by blocking a port, preventing it from receiving or forwarding data frames. An STP loop occurs when an STP blocking port erroneously transitions to the forwarding state. The loop protect feature checks to see whether BPDUs are not received on a nondesignated port, and then moves that port into the STP loop-inconsistent blocking state, instead of the learning or forwarding state.

[\[MPC3E MIC Overview\]](#)

- **Support for integrated routing and bridging (IRB) MAC synchronization in multichassis link aggregation for aggregated Ethernet (MX Series routers)**—MX Series routers with MPCs/MICs operating in multichassis link aggregation group (MC-LAG) with aggregated Ethernet configurations support integrated routing and bridging (IRB) MAC address synchronization. In earlier releases, VRRP was the only solution for sharing the same MAC across MC-LAG chassis for IRB interfaces. This feature is supported on 32-bit interfaces only and interoperates with earlier MPC/MIC releases.

[\[Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview\]](#)

- **Layer 2 Integrated routing and bridging functionality on MPC3E**—Junos OS Release 12.2 supports Layer 2 integrated routing and bridging (IRB) interfaces on MPC3E. IRB interfaces act as Layer 3 routing interfaces for bridge domains.
- **L2PT support on MPC3E**—Junos OS supports Layer 2 protocol tunneling (L2PT) on MX Series MPC3E interfaces.
- **BFD support for VCCV for Layer 2 VPNs, Layer 2 circuits, and VPLS on MPC3E (MX Series routers)**—Bidirectional Forwarding Detection (BFD) support for virtual circuit connectivity verification (VCCV) on MPC3E interfaces enables you to configure a control channel for a pseudowire, in addition to performing the corresponding OAM functions to be used over that control channel.

BFD provides a low-resource mechanism for the continuous monitoring of the pseudowire data path and for detecting data plane failures. This feature provides support for asynchronous mode BFD for VCCV as described in RFC 5885, *Bidirectional Forwarding Detection (BFD) for the Pseudowire Virtual Circuit Connectivity Verification (VCCV)*. Alternatively, you can use a ping operation to detect pseudowire failures. However, the processing resources required for a ping operation are greater than what is needed for BFD. In addition, BFD is capable of detecting data plane failure faster than a VCCV ping. BFD for pseudowires is supported for Layer 2 circuits (LDP-based), Layer 2 VPNs (BGP-based), and VPLS (LDP-based or BGP-based).

Starting with Release 12.2, Junos OS introduces a distributed model for BFD for VCCV. Previously BFD for VCCV followed a Routing Engine-based implementation, but as of Release 12.2 and later, BFD for VCCV follows a distributed implementation over PIC concentrators such as DPC, FPC, MPC, and MPC3E.

In Junos OS Release 12.2 and later, the periodic packet management process (ppmd) on the PIC concentrators handles the periodic packet management (send and receive) for BFD for VCCV. This enables Junos OS to create more BFD for VCCV sessions, and to reduce the time taken for error detection. Similarly, the distributed implementation improves the performance of Routing Engines because the Routing Engine resources used for BFD for VCCV implementation become available for Routing Engine-related applications when the BFD for VCCV-related processing moves to the PIC concentrators. The distributed BFD for VCCV implementation also enables the BFD for VCCV sessions to remain active across graceful restarts.

[[Configuring BFD for VCCV for Layer 2 VPNs, Layer 2 Circuits, and VPLS](#)]

## Multiprotocol Label Switching (MPLS)

- **Require BFD-triggered Packet Forwarding Engine local repair**—Enables you to configure BFD and MPLS ping for fast-failure detection without relying on fast physical level detection. With links between routers, when a route goes down, the rpd recalculates the next best path. When MPLS-FRR is enabled, ifl messages are flooded to all FPCs. The edge FPC enables the bypass MPLS LSP tunnel. Lastly, all routes are repaired and sent through the bypass MPLS LSP tunnel. The amount of time it takes to repair all routes is proportional to the number of routes.

[[BFD-Triggered Local Repair for Rapid Convergence](#)]

- **LDP downstream on demand**—The Label Distribution Protocol (LDP) is widely deployed in downstream unsolicited advertisement mode. As service providers integrate the access and aggregation networks into a single MPLS domain, LDP downstream on demand is needed to distribute the bindings between access and aggregation networks to keep the access node control plane as lightweight as possible and to avoid storing thousands of label bindings from upstream aggregation nodes. Instead of learning and storing all label bindings for all possible loopback addresses within the MPLS network, the access node uses LDP downstream on demand to request the label bindings for only the FECs corresponding to the loopback addresses of those egress nodes to which it has services configured.

To enable LDP downstream on demand on the router, include the **downstream-on-demand** statement at the [**edit protocols ldp session session-address**] hierarchy level. Specify the LDP downstream on demand policy using the **dod-request-policy** statement at the [**edit protocols ldp**] hierarchy level to send label bindings to the access node.

[[Example: Configuring LDP Downstream on Demand](#)]

- **Corouted bidirectional packet LSPs**—A corouted bidirectional packet LSP is a combination of two LSPs sharing the same path between a pair of ingress and egress nodes. It is established using the GMPLS extensions to RSVP-TE. This type of LSP can be used to carry any of the standard types of MPLS-based traffic, including Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs. You can configure a single BFD session for the bidirectional LSP (you do not need to configure a BFD session for each LSP in each direction). You can also configure a single standby bidirectional LSP to provide a backup for the primary bidirectional LSP.

Configure the **corouted-bidirectional** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level to specify that the LSP be established as a corouted bidirectional packet LSP. For the reverse path, configure the **corouted-bidirectional-passive** statement at the **[edit protocols mpls label-switched-path *lsp-name*]** hierarchy level to associate the LSP with the initial bidirectional LSP when it is signaled at the ingress router. You cannot configure both of these statements on the same LSP.

[\[Configuring Corouted Bidirectional LSPs\]](#)

- **Extends support for filtering MPLS-tagged IPv4 packets based on match conditions to T4000 Type 5 FPC (T4000-FPC5-3D)**—The support for filtering MPLS-tagged IPv4 packets based on IP parameters of up to five MPLS stacked labels is extended to the T4000 Type 5 FPC.

[\[Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic\]](#)

- **Support for filtering MPLS-tagged IPv6 packets based on match conditions on T4000 Type 5 FPC (T4000-FPC5-3D)**—Junos OS supports filtering MPLS-tagged IPv6 packets based on IP parameters of up to five MPLS stacked labels.

To configure the filter match conditions for the **mpls** family based on IP parameters, include the **from** statement at the **[edit firewall family *family-name* filter *filter-name* term *term-name*]** hierarchy level:

```
from {  
    match-conditions;  
}
```

[\[Standard Firewall Filter Match Conditions for MPLS-Tagged IPv4 or IPv6 Traffic\]](#)

- **Point-to-multipoint LSP traceroute support for T4000 routers**—You can use the label-switched path (LSP) trace utility to diagnose data plane failures in point-to-multipoint LSPs. To trace a point-to-multipoint LSP, issue the **traceroute mpls rsvp multipoint** command. The command also includes both the **egress** option, which enables you to specify a particular endpoint, and the **tth** option, which enables you to limit the number of hops to trace.

[\[traceroute mpls rsvp\]](#)

---

## Multicast

- **Controlling PIM resources for multicast VPNs**—Junos OS 12.2 introduces the following PIM configuration options to protect against potential misbehaving customer edge (CE) devices and VPN routing and forwarding (VRF) routing instances:
  - Limit the number of accepted PIM joins for any-source groups (\*G) and source-specific (S,G) groups. You can optionally configure a system log warning threshold, which allows you to generate and review system log messages to detect whether an excessive number of PIM joins have been received on the device. The system log warning threshold is a percentage of the configured PIM join limit. You can further specify a log interval, which is the amount of time (in seconds) between the log messages. To configure PIM join limits and the associated logging threshold,

include the **sglimit maximum *limit* <threshold value> <log-interval seconds>** statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* protocols pim]
  - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim]
  - [edit protocols pim]
  - [edit routing-instances *routing-instance-name* protocols pim]
- Limit the number of received PIM register messages on a rendezvous point (RP). You can optionally configure a system log warning threshold, which allows you to generate and review system log messages to detect whether an excessive number of PIM register messages have been received on the device. The system log warning threshold is a percentage of the configured PIM register message limit. You can further specify a log interval, which is the amount of time (in seconds) between the log messages. To configure PIM register message limits and the associated logging threshold, include the **register-limit maximum *limit* <threshold value> <log-interval seconds>** statement at the following hierarchy levels:
    - [edit logical-systems *logical-system-name* protocols pim rp]
    - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp]
    - [edit protocols pim rp]
    - [edit routing-instances *routing-instance-name* protocols pim rp]
  - Limit the number of group-to-RP mappings on an RP. You can optionally configure a system log warning threshold, which allows you to generate and review system log messages to detect whether an excessive number of group-to-RP mappings have been received on the device. The system log warning threshold is a percentage of the configured group-to-RP mapping limit. You can further specify a log interval, which is the amount of time (in seconds) between the log messages. To configure group-to-RP mapping limits and the associated logging threshold, include the **group-rp-mapping maximum *limit* <threshold value> <log-interval seconds>** statement at the following hierarchy levels:
    - [edit logical-systems *logical-system-name* protocols pim rp]
    - [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* protocols pim rp]
    - [edit protocols pim rp]
    - [edit routing-instances *routing-instance-name* protocols pim rp]



**NOTE:** The group-to-RP mappings limit does not apply to static RP or embedded RP configurations.

As a result of these PIM configuration options, the **show pim statistics instance *instance-name*** command has been updated to display the configured limits and currently accepted values for PIM join states and PIM register limits. The PIM register limit values are displayed on RPs configured for PIM register limits.

[[Example: Configuring PIM State Limits](#)]

- **Nonstop active routing PIM support for draft-rosen MVPNs**—Starting with Release 12.2, Junos OS extends the nonstop active routing PIM support to draft-rosen MVPNs. It enables nonstop active routing-enabled devices to preserve draft-rosen MVPN-related information—such as default and data MDT states—across switchovers. In releases earlier than Release 12.2, nonstop active routing PIM configuration was incompatible with draft-rosen MVPN configuration.

The backup Routing Engine sets up the default multicast distribution tree (MDT) based on the configuration and the information it receives from the master Routing Engine, and keeps updating the default MDT state information.

However, for data MDTs, the backup Routing Engine relies on the master Routing Engine to provide updates when data MDTs are created, updated, or deleted. The backup Routing Engine neither monitors data MDT flow rates nor triggers a data MDT switchover based on variations in flow rates. Similarly, the backup Routing Engine does not maintain the data MDT delay timer or timeout timer. It does not send MDT join TLV packets for the data MDTs until it takes over as the master Routing Engine. After the switchover, the new master Routing Engine starts sending MDT joins TLV packets for each data MDT, and it also resets the data MDT timers. Note that the expiration time for the timers might vary from the original values on the previous master Routing Engine.



**NOTE:** Nonstop active routing support for PIM does not include support for next-generation MVPNs. The commit fails if you configure nonstop active routing for PIM on devices configured for next-generation MVPN setups.

[[Nonstop Active Routing System Requirements](#)]

- **Support for PIM automatic make-before-break (MBB) join load balancing**—Ensures that PIM joins are evenly redistributed to all upstream PIM neighbors on an ECMP path. When an interface is added to an ECMP path, MBB provides a switchover to an alternate path with minimal traffic disruption. The feature can be enabled by using the **automatic** statement at the **[edit protocols pim join-load-balance]** hierarchy level. When a new neighbor is available, the time taken to create a path to the neighbor (standby path) can be configured by using the **standby-path-creation-delay seconds** statement at the **[edit protocols pim]** hierarchy level. In the absence of this statement, the standby path is created immediately and the joins are redistributed as soon as the new neighbor is added to the network. For a join to be moved to the standby path in the absence of traffic, the **idle-standby-path-switchover-delay seconds** statement is configured at the **[edit protocols pim]** hierarchy level. In the absence of this statement, the join is not moved until traffic is received on the standby path.



[[Example: Configuring PIM Make-Before-Break \(MBB\) Join Load Balancing](#)]

- **BFD client support for PIM IPv6**—Enables you to configure BFD liveness detection for IPv6 interfaces using Protocol Independent Multicast (PIM). Bidirectional Forwarding Detection (BFD) enables rapid detection of communication failures between adjacent systems. By default, authentication for BFD sessions is disabled. However, when you run BFD over Network Layer protocols, the risk of service attacks can be significant.

[[Example: Configuring BFD Liveness Detection for PIM IPv6](#)]

## Network Management and Monitoring

- **Updated MIB for IPv6 networks**—Junos OS Release 12.2 and later support the IP Forwarding MIB table and related objects used for forwarding IP packets in IPv6 networks (in addition to IPv4 networks), as per RFC 4292. The `inetCidrRouteTable` table displays IP version-independent multipath CIDR routes. The `inetCidrRouteNumber` object indicates the number of current routes in the `inetCidrRouteTable` table. The `inetCidrRouteDiscards` object counts the number of valid routes that are discarded from the `inetCidrRouteTable` table.

[[IP Forwarding MIB](#)]

- **Support for RFC 4087**—Junos OS Release 12.2 and later support two standard tables of the IP Tunnel MIB for managing tunnels of any type over IPv4 and IPv6 networks. The `tunnellfTable` table provides information about the tunnels known to a router. The `tunnellnetConfigTable` table displays information about the dynamic creation of tunnels, and mapping of endpoint addresses to the current interface index value.

[[Standard SNMP MIBs Supported by Junos OS](#)]

- **Junos OS support for proxy SNMP agent**—Junos OS enables you to assign one of the devices in the network as a proxy SNMP agent through which the network management system (NMS) can query other devices in the network. When you configure a proxy, you can specify the names of devices to be managed through the proxy SNMP agent.

When the NMS queries the proxy SNMP agent, the NMS specifies the community name (for SNMPv1 and SNMPv2) or the context and security name (for SNMPv3) associated with the device from which it requires the information.



**NOTE:** If you have configured authentication and privacy methods and passwords for SNMPv3, those parameters are also specified in the query for SNMPv3 information.

To configure a proxy SNMP agent and specify devices to be managed by the proxy SNMP agent, you can include the following configuration statements at the `[edit snmp]` hierarchy level:

```
proxy proxy-name {
  device-name device-name;
  <version-v1 | version-v2c> {
    snmp-community community-name;
    no-default-comm-to-v3-config;
  }
}
```

```

version-v3 {
    security-name security-name;
    context context-name;
}
logical-system logical-system {
    routing-instance routing-instance;
}
routing-instance routing-instance;
}

```

- The **proxy** statement enables you to specify a unique name for the proxy configuration.
- The **version-v1**, **version-v2**, and **version-v3** statements enable you to specify the SNMP version.
- The **no-default-comm-to-v3-config** statement is an optional statement at the **[edit snmp proxy proxy-name <version-v1 | version-v2>]** hierarchy level that when included in the configuration requires you to manually configure the statements at the **[edit snmp v3 snmp-community community-name]** and **[edit snmp v3 vacm]** hierarchy levels.  
  
If the **no-default-comm-to-v3-config** statement is not included at the **[edit snmp proxy proxy-name <version-v1 | version-v2>]** hierarchy level, the **[edit snmp v3 snmp-community community-name]** and **[edit snmp v3 vacm]** hierarchy level configurations are automatically initialized.
- The **logical-system** and **routing-instance** statements are optional statements that enable you to specify logical system and routing instance names if you want to create proxies for logical systems or routing instances on the device.



**NOTE:** The community and security configuration for the proxy should match the corresponding configuration on the device that is to be managed.



**NOTE:** Because the proxy SNMP agent does not have trap forwarding capabilities, the devices that are managed by the proxy SNMP agent send the traps directly to the network management system.

You can use the **show snmp proxy** operational mode command to view proxy details on a device. The **show snmp proxy** command returns the proxy names, device names, SNMP version, community/security, and context information. [*Network Management*]

## Routing Policy and Firewall Filters

- **Extends filter and policer feature support to T4000 Type 5 FPC (T4000-FPC5-3D)**—The following filter and policer features supported on the T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES) are also supported on the T4000 Type 5 FPC (T4000-FPC5-3D):
  - Label-switched path (LSP) policers
  - Address Resolution Protocol (ARP) policers

- Tricolor marking policers
- Forwarding table filters
- Filter-based forwarding
- Prefix-specific actions

The following filter and policer features supported on the T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES) are not supported on the T4000 Type 5 FPC (T4000-FPC5-3D):

- Service PIC-related filters.
- Applying a policer at the logical interface level.
- Filter actions such as **ipsec-sa**, **service-accounting**, and **service-filter-hit**.
- The **dscp 0** action is not supported during the interoperation between a T1600 Enhanced Scaling Type 4 FPC and a T4000 Type 5 FPC.
- Shared bandwidth policer.
- A filter attached at the Layer 2 application point (that is, at the logical interface level) is unable to match with the forwarding class of a packet that is set by a Layer 3 classifier such as DSCP, DSCP V6, **inet-precedence**, or **mpls-exp**.
- Using **interface-group** and **interface-group-except** as match conditions for the VPLS family filter.
- Applying filters at **set interfaces lo0 unit 0 family any filter input filter-name**.
- For a three-color policer operating in color-aware mode and when the PLP of the input packet is medium-low, the color of the input packet to the policer is mapped to the color yellow.

In such a scenario, if the color of the input packet remains unchanged, the policer operates in the following way:

- On a T1600 Enhanced Scaling Type 4 FPC (T1600-FPC4-ES), the PLP of the output packet remains medium-low.
- On a T4000 Type 5 FPC (T4000-FPC5-3D), the PLP of the output packet is marked as medium-high.

Because of this difference, for any applications (such as rewrite and WRED selection on egress interface) that use PLP, the packets are treated differently for the same flow depending on the FPC type (T1600 Enhanced Scaling FPC4 (T1600-FPC4-ES) or T4000 FPC5 (T4000-FPC5-3D)) on which the policer is applied.

*[[Configuring Policers for LSPs](#), [Three-Color Policer Configuration Overview](#), [Configuring Forwarding Table Filters](#), [Filter-Based Forwarding Overview](#), [Prefix-Specific Action Configuration](#), [Color-Aware Mode](#)]*

- **Hierarchical policer support for T4000 Type 5 FPC (T4000-FPC5-3D)**—Type 5 FPCs on T4000 routers support hierarchical policers only at the interface family level.



**NOTE:** Support for hierarchical policers at the physical and logical interface levels requires the presence of an IQE PIC. Because the T4000 Type 5 FPC does not have an IQE PIC, hierarchical policers are not supported at the physical and logical interface levels on this FPC.

---

*[[hierarchical-policer, aggregate \(Hierarchical Policers\)](#)]*

- **Change to routing policy match condition**—Starting in Junos OS Release 12.2 the routing policy match condition **from protocol rtarget** has been changed to **from protocol route-target**. The **from protocol rtarget** match condition is hidden and continues to work in Junos OS Release 12.2 and later. However, if you configure **from protocol route-target** and then downgrade to Junos OS Release 12.1 or earlier, the configuration will not commit.
- **Filter-based forwarding to a specific outgoing interface or destination IP address (MX Series routers with MPCs)**—Enables you to use filter-based forwarding (sometimes also referred to as policy-based routing or PBR) to apply a match condition and send packets to a certain outgoing interface or to a certain IPv4 or IPv6 address. To configure, use the **next-interface**, **next-ip**, or **next-ip6** firewall filter action.

*[[Example: Configuring Filter-Based Forwarding to a Specific Outgoing Interface or Destination IP Address](#)]*

## Routing Protocols

---

- **Origin validation for BGP**—Enables BGP to recognize when an autonomous system (AS) begins advertising all or part of another company's assigned network. BGP recognizes the error and responds in a way that avoids service interruptions. To configure, include the **validation** statement (and associated child statements) at the **[edit routing-options]** hierarchy level. Also configure a policy with the **from validation-database** match condition, the **then validation-state** action, and the extended community (origin validation state).

*[[Example: Configuring Origin Validation for BGP](#)]*

## Services Applications

---

- Starting with Junos OS Release 11.4R11, interim-logging is supported with NAT64 on microkernel (MS-DPC) platforms. The configuration statement **pba-interim-logging-interval** under the **[interfaces services-options]** hierarchy level enables the feature for NAT64.

## Subscriber Access Management

---



**NOTE:** Although present in the code, the subscriber management features are not supported in Junos OS Release 12.2R9. Documentation for subscriber management features is included in the Junos OS Release 12.2 documentation set.

---

- **Junos OS subscriber management scaling values (M120, M320, and MX Series routers)**—A spreadsheet is available online that lists scaling values supported for Junos OS subscriber management beginning with Junos OS Release 10.1. Access the *Subscriber Management Scaling Values (XLS)* spreadsheet from the Downloads box at [http://www.juniper.net/techpubs/en\\_US/junos12.2/information-products/pathway-pages/subscriber-access/index.html](http://www.juniper.net/techpubs/en_US/junos12.2/information-products/pathway-pages/subscriber-access/index.html). You can also substitute the number of the latest Junos OS release for the *12.2 release-number*. For example, *...en\_us/junos11.1/...*

[Subscriber Management Scaling]

- **Scaling enhancements**—This release enables significant scaling and performance gains applicable to a broad range of broadband edge deployment models. Absolute scaling and performance numbers achievable are influenced by a number of factors including deployment model, software configuration, and hardware configuration. *Applicable subscriber scaling licenses apply.*

Maximum scaling and performance for broadband edge configurations require the RE-S-1800 Routing Engine and MPC2 access-facing modules.

[Subscriber Management Scaling]

- **Scaling resource management**—The memory resource management feature enables additional system protection by limiting subscriber logins during times of high memory utilization. Limiting subscriber logins helps avoid resource exhaustion. As utilization decreases, the subscriber limits are removed.

[Subscriber Management Scaling]

- **Support for configuring dynamic VLAN subscriber interfaces using agent-circuit-identifier information (MX Series routers with MPCs/MICs)**—Enables you to configure dynamic VLAN subscriber interfaces for DHCP and PPPoE subscribers based on agent-circuit-identifier information. To use this feature, you must configure the dynamic VLAN subscriber interfaces on MPC/MIC modules that face the access side of the network in an MX Series router.

In Ethernet-based subscriber access networks, DHCP and PPPoE subscribers are uniquely identified either by means of VLAN encapsulation (that is, the S-VLAN ID tag and the VLAN ID tag), or by insertion of the agent-circuit-identifier string in DHCP and PPPoE control messages.

For dynamic VLAN subscriber interfaces with single-tagged, untagged, or double-tagged VLAN encapsulation, you can configure the router to examine the DHCP and PPPoE control packets to extract the agent-circuit-identifier information in order to build a unique VLAN subscriber interface. The agent-circuit-identifier value is a string that uniquely identifies the subscriber's access node and the DSL line on the access node. For DHCP traffic, the agent-circuit-identifier string is in the DHCP option 82 field of DHCP messages. For PPPoE traffic, the agent-circuit-identifier string is in the DSL Forum Agent-Circuit-ID VSA [26-1] of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets.

Configuring dynamic VLAN subscriber interfaces based on agent-circuit-identifier information is particularly useful in configurations with multiple DHCP and PPPoE subscriber sessions per household. Because DHCP and PPPoE traffic sent to the router from the same household has the same agent-circuit-identifier information, the router

groups these DHCP and PPPoE subscribers in the same agent-circuit-identifier interface set. An *agent-circuit-identifier interface set* is a logical collection of subscriber interfaces that originate at the same household or on the same access-loop port. Grouping subscribers into agent-circuit-identifier interface sets facilitates application of subscriber-based services, such as class of service (CoS) and interface-shared filters, to all of the subscriber's interfaces.

Configuring a dynamic VLAN subscriber interface based on agent-circuit-identifier information involves the following basic steps:

1. Create a dynamic profile that defines the agent-circuit-identifier interface set.  
To reference the interface set, include the **interface-set** statement with the **\$junos-interface-set-name** predefined variable at the **[edit dynamic-profiles profile-name interfaces]** hierarchy level.
2. (Optional) Include attributes for PPPoE, CoS, and interface-shared filters in the dynamic profile for the agent-circuit-identifier interface set.  
For dynamic PPPoE subscriber interfaces, you can include the **max-sessions** statement at the **[edit dynamic-profiles profile-name interfaces interface-set "\$junos-interface-set-name" pppoe-underlying-options]** hierarchy level.
3. Configure the underlying VLAN interface to enable dynamic subscriber interface creation based on agent-circuit-identifier information.
  - For a statically created underlying VLAN interface, include the **auto-configure** stanza at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level.
  - For a dynamically created underlying VLAN interface, include the **auto-configure** stanza at the **[edit dynamic-profiles profile-name interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]** hierarchy level.
4. Associate the dynamic agent-circuit-identifier interface set with the logical subscriber interface.
  - In a dynamic profile for a PPPoE logical subscriber interface, include the **interface-set \$junos-interface-set-name interface pp0 unit \$junos-interface-unit** statement at the **[edit dynamic-profiles profile-name interfaces]** hierarchy level.
  - In a dynamic profile for an IP demultiplexing (IP demux) logical subscriber interface for DHCP subscribers, include the **interface-set \$junos-interface-set-name interface demux0 unit \$junos-interface-unit** statement at the **[edit dynamic-profiles profile-name interfaces]** hierarchy level.

To verify and manage dynamic VLAN configurations based on agent-circuit-identifier information, you can use the following operational commands:

- **clear auto-configuration interfaces interface-set**
- **show subscribers aci-interface-set-name**
- **show subscribers agent-circuit-identifier**

In addition, the output of the following operational commands has been enhanced to help you verify and manage this feature for DHCP and PPPoE subscribers:

- **show dhcp server binding detail**
- **show interfaces**
- **show pppoe interfaces**
- **show pppoe underlying-interfaces**
- **show subscribers detail**

[[Configuring Dynamic VLANs Based on Agent Circuit Identifier Information](#)]

- **Support for subscriber services over ATM networks (MX Series routers with MPCs and ATM MICs with SFP)**—By using the ATM MIC with SFP (model number MIC-3D-80C3-20C12-ATM) and a supported MPC, you can configure an MX Series router to support the following configurations that enable subscribers to access an MX Series router over an ATM network:

- PPPoE-over-ATM

PPP-over-Ethernet-over-ATM (PPPoE-over-ATM) configurations support both statically created and dynamically created PPPoE (**pp0**) logical subscriber interfaces over static ATM underlying interfaces. Most PPPoE and subscriber services features supported on terminated connections and tunneled (L2TP access concentrator, or LAC) connections are also supported for access to an MX Series router over an ATM network. You can dynamically apply subscriber services such as class of service (CoS) and firewall filters to the **pp0** logical subscriber interface by configuring the services in the dynamic profile that creates the static or dynamic **pp0** logical interface.

For PPPoE-over-ATM configurations on an MX Series router, you must configure the ATM underlying interface with PPPoE-over-ATM logical link control (LLC) encapsulation. To do so, include the **encapsulation ppp-over-ether-over-atm-llc** statement at the **[edit interfaces interface-name unit logical-unit-number]** hierarchy level.

You must configure the router to act as a PPPoE server (also known as a *remote access concentrator*) in PPPoE-over-ATM configurations on MX Series routers. Configuring the router to act as a PPPoE client in these configurations is not supported.

- Routed IP-over-ATM

Routed IP-over-ATM (IPoA) configurations support statically created IPv4 and IPv6 logical subscriber interfaces over static ATM underlying interfaces. (Dynamic creation of IPv4 or IPv6 interfaces is not supported.) Subscriber services such as CoS and firewall filters must also be statically configured; you cannot use a dynamic profile for this purpose.

Routed IPoA configurations on an MX Series router support two types of encapsulation on the ATM underlying interface:

- To configure routed IPoA encapsulation that uses LLC, you must configure the ATM underlying interface with ATM subnetwork attachment point (SNAP)

encapsulation. To do so, include the **encapsulation atm-snap** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

- To configure routed IPoA encapsulation that uses virtual circuit (VC) multiplexing, you must configure the ATM underlying interface with ATM VC multiplex encapsulation. To do so, include the **encapsulation atm-vc-mux** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

- Bridged IP-over-Ethernet-over-ATM

Bridged IP-over-Ethernet-over-ATM configurations support statically created IPv4 and IPv6 logical subscriber interfaces over static Ethernet interfaces over static ATM underlying interfaces. (Dynamic creation of IPv4, IPv6, or Ethernet interfaces is not supported.) Subscriber services such as CoS and firewall filters must also be statically configured; you cannot use a dynamic profile for this purpose.

For bridged IP-over-Ethernet-over-ATM configurations on an MX Series router, you must configure the ATM underlying interface with Ethernet-over-ATM LLC encapsulation. To do so, include the **encapsulation ether-over-atm-llc** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

- PPP-over-ATM

PPP-over-ATM (PPPoA) configurations support statically created PPP logical subscriber interfaces over static ATM underlying interfaces. (Dynamic creation of the PPP interfaces is not supported.) Most features supported for PPPoE configurations are also supported for PPP access to an MX Series router over an ATM network. You can dynamically apply subscriber services such as CoS and firewall filters to the static PPP logical subscriber interface by configuring the services in the dynamic profile that creates the PPP logical interface.

PPPoA configurations on an MX Series router support two types of encapsulation on the ATM underlying interface:

- To configure PPPoA encapsulation that uses LLC, you must configure the ATM underlying interface with PPP-over-AAL5 LLC encapsulation. To do so, include the **encapsulation atm-ppp-llc** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.
- To configure PPPoA encapsulation that uses VC multiplexing, you must configure the ATM underlying interface with PPP-over-ATM AAL5 multiplex encapsulation. To do so, include the **encapsulation atm-ppp-vc-mux** statement at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level.

Using these configurations enables the delivery of subscriber-based services, such as CoS and firewall filters, for subscribers accessing the router over an ATM network. In addition, PPPoE-over-ATM support on an MX Series router enables you to configure the router to dynamically create PPPoE logical subscriber interfaces over static ATM underlying interfaces only when needed; that is, when a subscriber logs in on the associated underlying interface. (Dynamic PPPoE over static ATM configurations are not supported on M Series routers and T Series routers.)

You can use the same basic statements, commands, and procedures to create, verify, and manage PPPoE-over-ATM, IPoA, IP-over-Ethernet-over-ATM, and PPPoA



configurations as the statements, commands, and procedures you use for static configurations on M Series routers and T Series routers, and for dynamic PPPoE configurations on MX Series routers.

[[Junos OS Subscriber Access Configuration Guide](#), [Junos OS ATM Interfaces Configuration Guide](#), [Junos OS Circuit Emulation Interfaces Configuration Guide](#)]

- **Sharing schedulers and scheduler maps across dynamic CoS**—The system generates unique identifiers (IDs) in dynamic profiles created for services. The generated unique IDs enable you to identify and configure separate parameter values with the same variable name. When applied to CoS, you can configure scheduler and scheduler map sharing. In client profiles, schedulers and scheduler maps must use the unique ID format. If the client profile uses the unique ID format and you want to have either scheduler or scheduler map sharing for service activation, you must configure the service profile in unique ID format. Generating unique IDs based on schedulers and scheduler maps eliminates duplication and improves router performance and scalability.

To enable this feature, include the variables for CoS in the client or service dynamic profile. For example, to have scheduler map and scheduler sharing, you need to define variables for the scheduler maps **smap\_data** and **smap\_voice** and for the schedulers **data\_sched** and **voice\_sched** in the dynamic profile. You then add the scheduler maps and schedulers in the variable format such as **\$smap\_data**, **\$smap\_voice**, **\$data\_sched**, and **\$voice\_sched**, respectively to the **class-of-service** hierarchy.

[[Access Profiles and Service Profiles Overview](#)]

- **ANCP enhancements for VLAN demux over aggregated Ethernet interfaces and RADIUS (MX Series routers)**—ANCP can perform class-of-service traffic shaping for a household, for individual PPPoE sessions within a household, or for both. ANCP supports VLAN demux over aggregated Ethernet interfaces, with or without interface sets. This support includes the following:
  - Mapping agent circuit identifiers (ACIs) to interfaces.
  - Dynamically updating the CoS process with the adjusted downstream data rate. ANCP receives the actual data rate from the access node and then adjusts it according to its configuration before updating CoS. (The upstream rate is not provided to CoS because it does not shape upstream traffic.)
  - OAM support for managing PPPoE sessions.

ANCP can now make CoS-related adjustments to the upstream data rate it receives from the access node. (In earlier releases, only downstream rate adjustment was possible.) ANCP can report both the adjusted and the unadjusted values to authd for RADIUS authentication and accounting. Rate adjustment and rate reporting are supported on the following interfaces, with or without interface sets:

- VLAN over Ethernet
- VLAN demux over aggregated Ethernet

By reporting adjusted data rates, ANCP enables RADIUS to allocate the appropriate services (including class of service) to PPPoE sessions during authentication. The

reports also enable RADIUS accounting to track the actual class of service provided for PPPoE sessions, which in turn enables accurate billing for subscriber services.

ANCP stores the DSL attributes that it receives from access nodes in the shared database. The ANCP DSL attributes are mapped by authd to the Juniper Networks DSL VSAs used by RADIUS. RADIUS uses these attributes during authentication and accounting for PPPoE sessions on the subscriber access line. The attributes persist even when the ANCP session to a given node has ended, enabling RADIUS to later apply these attributes to new sessions on that subscriber access line. To remove the attributes, you must delete the access line from the ANCP configuration.

The RADIUS profile must be configured to include the **juniper-dsl-attributes** option, or authd does not report the attributes to RADIUS. If the ANCP DSL attributes are unavailable, the session's advisory upstream and downstream data rates are mapped to the calculated upstream and downstream data rate VSAs. These VSAs alone are then provided to RADIUS.

For successful authentication and accounting by RADIUS, AAA has to correlate PPPoE sessions with their access lines and their associated DSL attributes. Some access nodes provide the ACI in PADI/PADR packets for the PPPoE sessions.

When the ACI is not provided in a 1:1 VLAN model with interface sets, you must associate the underlying interface for the sessions with the ACI and the interface set. If you do not configure this association, then only the advisory traffic rates are provided to RADIUS. This configuration has no effect when the ACI is provided by the access node.

For the N:1 VLAN model with interface sets, the access node must provide the ACI. If you configure the underlying interface for this model when the access node does not provide the ACI, PPPoE sessions could be incorrectly correlated with access lines.

To map an ACI to a static VLAN demux interface, include the **access-identifier identifier** statement, and optionally the **neighbor neighbor-ip-address** statement, at the **[edit protocols ancp interfaces demux0.logical-unit-number]** hierarchy level.

To configure advisory upstream and downstream data rates on a static VLAN demux interface, include the **upstream-rate rate** or **downstream-rate rate** statements at the **[edit interfaces demux0 unit logical-unit-number]** hierarchy level.

To configure an underlying interface for the PPPoE sessions in an interface set, include the **underlying-interface interface-name** statement at the **edit protocols ancp interfaces interface-set interface-set-name** hierarchy level.

[\[ANCP Operations in Different Network Configurations\]](#)

- **PPP options and keepalives supported for L2TP LNS subscribers per interface (MX Series routers)**—You can configure PPP options for LNS subscribers on inline services (si) interfaces on a per-interface basis. In earlier releases, you applied a configuration for PPP options only with a user group profile, which specifies the same configuration for all subscribers processed through a particular LAC client. The new support matches the existing behavior for terminated PPPoE subscribers on pp0 interfaces and uses the following existing statements:

```
ppp-options {  
  chap;  
  pap;  
}
```

```
}
```

For dynamically created si interfaces, include the statements at the **[edit dynamic-profiles *profile-name* interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]** hierarchy level.

For statically configured si interfaces, include the statements at the **[edit interfaces *si-slot/pic/port* unit *logical-unit-number*]** hierarchy level.



**BEST PRACTICE:** Although all other statements subordinate to **ppp-options**—including those subordinate to **chap** and **pap**—are supported, they are typically not used for subscriber management. We recommend that you leave these other statements at their default values.

Similarly, you can configure PPP keepalives on a per-interface basis, whereas in earlier releases you configured PPP keepalives only with a user group profile.

For dynamic si interfaces, include the **keepalives** statement at the **[edit dynamic-profiles *profile-name* interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit"]** hierarchy level.

For static si interfaces, include the **keepalives** statement at the **[edit interfaces *si-slot/pic/port* unit *logical-unit-number*]** hierarchy level.

When you change the PPP keepalive configuration in a user group profile, the modified configuration affects only new sessions logging in. Sessions that exist at the time of the change are not affected.

When you configure the PPP options or PPP keepalives for L2TP LNS subscribers both on the si interface and in user group profiles, the inline service interface configuration takes precedence over the group profile configuration.

[\[Applying PPP Attributes to L2TP LNS Subscribers Per Inline Service Interface\]](#)

- **Support for CoS on dynamic VLAN subscriber interfaces using agent-circuit-identifier information (MX Series routers with MPCs/MICs)**—Enables you to configure specified class-of-service (CoS) attributes using a dynamic interface set. Because the interface sets corresponding to VLANs using agent-circuit-identifier (ACI) information are created dynamically, you can apply CoS attributes, such as shaping, at the household level. You must set and define the CoS policy for the ACI virtual VLAN interface set using the ACI set profile (not the subscriber profile). CoS on dynamic VLANs includes support for level 3 or level 2 scheduler nodes for a dynamic interface set. You can also configure a traffic control profile and a remaining traffic control profile for a dynamic interface set. CoS on dynamic VLANs enables you to configure a dynamic scheduler map for a traffic control profile that is used by a dynamic interface set. In this case, the dynamic scheduler map must use the UID format. This feature ensures that a subscriber receives a minimum bandwidth (guaranteed rate) and a maximum bandwidth (shaping rate), which reduces network operational expenses by providing centralized management of the network.

To enable this feature, include attributes for CoS in the dynamic profile for the agent-circuit-identifier interface set. For example, for dynamic CoS subscriber interfaces,

you can include the **output-traffic-control-profile** statement or **output-traffic-control-profile-remaining** statement at the **[edit dynamic-profiles profile-name class-of-service interfaces interface-set "\$junos-interface-set-name"]** hierarchy level.

[ [CoS for Subscriber Access Overview](#) ]

- **Support for dynamic interface-shared filters (MX Series routers with MPCs/MICs)**—Enables you to configure a type of dynamic filter attachment. Interface-shared filters can be defined statically or dynamically, but can only be applied using dynamic profiles, and are supported for both client and service sessions. The same interface-shared instance can be attached to multiple interfaces only if these interfaces reference the same interface-shared filter name and have the same shared-name. The shared-name can either be populated from **\$junos-interface-set-name**, where the value comes from the related client session, or a service session variable.

With VLAN subscriber interfaces that use the agent-circuit-identifier information, many subscribers share the same underlying logical interface. Because some of these subscribers are related to each other as part of the same household, you must apply an interface-shared filter to the subscriber logical interfaces that make up the household to be able to filter and police these related subscribers at a household level. All interfaces that share the same interface-shared filter instance share the same set of counters and policer actions.

The base filter name of a parameterized filter is assigned depending upon the profile name and the contents of the filter definition. Therefore, when interface-shared filter is used with parameterized filters, all service sessions expecting to share the same instance of an interface-shared filter must have the exact same parameterized filter and profile. A service session should expect a different instance of the interface-shared filter if either the parameterized filter or the profile is different.

To use this feature, you must configure the filter by using the **interface-shared** statement at the **[edit firewall family [inet|inet6] filter filter-name]** hierarchy level. To attach this type of filter to an interface, the shared-name must be defined, and the interface must be a dynamic interface, which is defined within a dynamic-profiles hierarchy.

[ [Interface-Shared Filters Overview](#) ]

- **Option 82 suboptions in authentication usernames for autosense VLANs (MX Series routers)**—You can specify the option 82 suboptions that are concatenated with the username during the authentication process for autosense VLANs. The option 82 value used in creating the username is based on the option 82 value that is encoded in the incoming DHCP discover packet. You can specify either both or neither of the Agent Circuit ID (suboption 1) and the Agent Remote ID (suboption 2). If you specify both, the Agent Circuit ID is supplied first, followed by a delimiter, and then the Agent Remote ID. If you specify that neither suboption is supplied, the raw payload of option 82 from the PDU is concatenated to the username. The use of option 82 suboptions is supported for DHCPv4 discover packets only.

Use the **option-82 circuit-id remote-id** statement at the **[edit interfaces interface-name auto-configure vlan-ranges authentication username-include]** hierarchy level to configure option 82 support for autosense VLANs.

[ [Option 82 Suboptions in Authentication Usernames for Autosense VLANs](#) ]

- **S-VLAN-based shaping support for dynamic profiles**—This release supports CoS traffic shaping of service VLAN (S-VLAN) interface sets in dynamic profiles. An interface set enables you to group interfaces into a logical group and provide the same level of service for that group of subscribers.

This feature requires the introduction of the following internal dynamic variables:

- `$junos-svlan-interface-set-name`—Locally generated interface set name for use by dual-tagged VLAN interfaces based on the outer tag of the dual-tagged VLAN. The format of the generated variable is *physical\_interface\_name - outer\_VLAN\_tag*.
- `$junos-tagged-vlan-interface-set-name`—Locally generated interface set name used for grouping logical interfaces stacked over logical stacked VLAN demux interfaces for either a 1:1 (dual-tagged; individual client) VLAN or N:1 (single tagged; service) VLAN. The format of the generated variable differs with VLAN type. For dual-tagged (client) VLANs, the format of the generated variable is *physical\_interface\_name - outer\_VLAN\_tag - inner\_VLAN\_tag*. For single tagged (service) VLAN, the format of the generated variable is *physical\_interface\_name - VLAN\_tag*.

To configure VLAN-based shaping, include the **interface-set** statement, along with the desired dynamic variable, at the **[edit dynamic-profiles dynamic-profile-name interfaces]** hierarchy level. You must also include the **interface** statement, along with the **demux0** interface, at the **[edit dynamic-profiles dynamic-profile-name interfaces interface-set dynamic-variable]** hierarchy level, and the **unit** statement, along with the *\$junos-interface-unit* dynamic variable for the dynamically created set units, at the **[edit dynamic-profiles dynamic-profile-name interfaces interface-set dynamic-variable interface demux0]** hierarchy level.

In addition to configuring the interface set for the dynamic profile, you must also include the expected interface set name for each physical or aggregated interface that you want to be part of the interface set. For example, to specify the expected interface set name for aggregated Ethernet interface ae0 and outer VLAN tag 111, include **ae0-111** at the **[edit class-of-service interfaces demux0]** hierarchy level.

[ [CoS for Interface Sets of Subscribers Overview](#) ]

## User Interface and Configuration

---

- **Support for configuring CLI breadcrumbs**—The output of the **show configuration** operational mode command and the **show** configuration mode commands can be configured to display configuration breadcrumbs that indicate the exact location in the hierarchy of the output being viewed. To enable the feature, configure the **configuration-breadcrumbs** statement at the **[edit system login class *class-name*]** hierarchy level.

[*Example: Enabling Configuration Breadcrumbs*]

## VPNs

---

- **VPLS improved convergence time for multihomed sites**—You can improve the convergence time for VPLS multihomed sites by configuring the **best-site** statement at the **[edit routing-instances *routing-instance-name* protocols vpls site *site-name*]** hierarchy level and the **mac-flush** statement at the **[edit routing-instances *routing-instance-name* protocols vpls]** hierarchy level. The **best-site** statement is new for Junos OS Release 12.2 and designates the site as the most preferable site for the provider edge (PE) router. The **mac-flush** statement is an existing statement. It enables media access control (MAC) flush processing for the VPLS routing instance or for the mesh group under a VPLS routing instance. MAC flush processing removes MAC addresses from the MAC address database that have been learned dynamically. With the dynamically learned MAC addresses removed, MAC address convergence requires less time to complete.

[*Example: VPLS Multihoming, Improved Convergence Time*]

- **Layer 3 VPN localization**—Layer 3 VPN localization provides a mechanism for localizing routes of instance type **vrf** or **virtual-routers** to specific Packet Forwarding Engines to help maximize the number of routes or VRFs that a router can handle.

To accomplish this, the Layer 3 VPN routes are installed only on the CE-facing Packet Forwarding Engine. By doing this, you can optimize the Packet Forwarding Engine memory. By Layer 3 VPN localization, the number of VPN IP routes that can be handled can be increased by using multiple Layer 3 VPN instances that are distributed across multiple Packet Forwarding Engines.

You can use the following statements at the **[edit routing-instances *routing-instance-name* routing-options]** hierarchy level to configure route localization for VRF:

- **localize**—Include this statement to localize routing-instance routes to a specific Packet Forwarding Engine hardware. This statement is applicable to **inet** and **inet6** families in the routing instance. It is not applicable for address families such as ISO and MPLS.

For routing instances of type **vrf**, the **localize** statement can be specified along with the **vrf-table-label**. You can also configure the statement in a VRF table that includes a **vt-** interface. If both **localize** and **vrf-table-label** are specified, the **localize** statement takes precedence for an L3VPN route label allocation. Similarly, if a **vt-** interface is

configured along with the **localize** statement, the **vt-** interface takes precedence for an L3VPN route label allocation.

You can configure the following options for this statement:

- **unicast-only**—Localizes unicast routes for the route tables associated with the routing instance. If the **localize** statement is configured without this option, the device localizes both unicast and multicast routes for the route tables associated with the routing instance.
- **source-class-usage**—Enables the Packet Forwarding Engine for source-prefix lookup in the context of a per-Packet Forwarding Engine table next hop at the egress CE-facing Packet Forwarding Engine. Include this statement for a VRF routing instance for packets coming from the MPLS core.

To enable flexible label allocation for localization, you can specify a different label allocation policy when you configure a VRF with localization. Use the **per-table-localize** option for the **label-allocation** statement at the **[edit policy-options policy-statement policy-statement-name term term-name then]** hierarchy level.

The **per-table-localize** label allocation policy is only applicable if the VRF is configured with the **localize** statement.

Issue the **show route instance detail** command to view VRF route localization details. You can also use the **show route table mpls protocol vpn** command to view details of the VPN route next hops.

*[Example: Configuring Layer 3 VPN Localization]*

- **Egress protection for Layer 3 VPN edge protection**—Typically, Layer 3 VPN service-restoration for multihomed customer edge (CE) routers depends on the ingress provider edge (PE) router to detect the egress PE link or node failure and switch traffic to the backup PE router. To achieve faster restoration, you can use a protector mechanism for the PE router to perform local restoration of the service immediately in case of an egress PE node failure. This mechanism is known as egress protection and requires the router at the point of local repair (PLR) router to redirect VPN traffic to a protector PE router for fast reroute of traffic. When you configure egress protection, the PLR detects the protected PE link or node failure and reroutes traffic through the protector PE router using the backup LDP-signaled LSP. The PLR uses per-prefix LFAs to program the backup next hop through the protected PE router, and traffic is forwarded to the CE routers using the alternate paths. This restoration is done quickly after the PLR router detects the PE egress node or link failure.

You can use the following configuration statements to configure egress protection:

- **egress-protection**—Configures protector information for the Layer 3 VPN and edge protection virtual circuit for the MPLS protocol. It also configures the context identifier at the **[edit protocols mpls]** hierarchy level.

The **egress-protection** statement configured as **unicast** at the **[edit protocols bgp group group-name family inet-vpn]**, **[edit protocols bgp group group-name inet6-vpn]**, or **[edit protocols bgp group group-name iso-vpn]** hierarchy levels contains a context ID for the context identifier. Include this statement to enable egress protection for the

configured BGP VPN network layer reachability information (NRLI). This configuration is required only on the protected PE and is not on the protector router.

The **egress-protection** statement configured at the **[edit routing-instances]** hierarchy level holds the context identifier of the protected PE. Include this statement in the configuration only on the primary PE router for outbound BGP updates for the next hops.

- **context-identifier**—Specifies an IPv4 address used to define the pair of PE routers participating in the egress protection LSP. The context identifier is used to assign an identifier to the protector PE router. The identifier is propagated to the other PE routers participating in the network, making it possible for the protected egress PE router to signal the egress protection LSP to the protector PE router.

[ [Example: Configuring MPLS Egress Protection for Layer 3 VPN Services](#) ]

- **Support for configuring more than one million Layer 3 VPN labels**—For Layer 3 VPNs configured on Juniper Networks routers, Junos OS normally allocates one inner VPN label for each customer edge (CE)-facing virtual routing and forwarding (VRF) interface of a provider edge (PE) router. However, other vendors allocate one VPN label for each route learned over the CE-facing interfaces of a PE router. This practice increases the number of VPN labels exponentially, which leads to slow system processing and slow convergence time.

For Juniper Networks routers participating in a mixed vendor network with more than one million Layer 3 VPN labels, include the **extended-space** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress l3vpn]** hierarchy level. The **extended-space** statement is disabled by default.

We recommend that you configure the **extended-space** statement in mixed vendor networks containing more than one million BGP routes to support Layer 3 VPNs. However, because using this statements can also enhance the Layer 3 VPN performance of Juniper Networks routers in networks where only Juniper Networks routers are deployed, we recommend configuring the statement in these networks as well.

[ [Accepting BGP Updates with Unique Inner VPN Labels in Layer 3 VPNs](#) ]

- **Proxy BGP route target filtering**—This feature (also known as proxy route target constrain, or proxy RTC) permits the generation of route target membership (RT membership) for devices that do not support route target filtering. This eases the deployment of route target filtering in networks where it is incompletely deployed or not fully supported. Proxy BGP route target filtering allows you to distribute proxy RT membership advertisements created from the received BGP VPN routes to other devices in the network that need them. These are known as proxy advertisements because the device creates the RT membership on behalf of its peers without the route target filtering functionality. Proxy BGP route target filtering uses BGP route target extended communities that are exported to a specific BGP speaker to generate the route targets. Generated proxy RTC routes are stored in the **bgp.rtarget.0** routing table.



To configure proxy BGP route target filtering, include the **proxy-generate** `<route-target-policy route-target-policy-name>` statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* protocols bgp group *group-name* family route-target]
- [edit logical-systems *logical-system-name* protocols bgp group *group-name* neighbor *address* family route-target]
- [edit protocols bgp group *group-name* family route-target]
- [edit protocols bgp group *group-name* neighbor *address* family route-target]

You can also configure a policy to further control route target filtering routes. This functionality applies to both BGP route target filtering and proxy BGP route target filtering. You define a list of route target prefixes to use in a routing policy and then apply those route target prefixes to the routing policy.

To define a list of route target prefixes to use in a routing policy, include the **rtf-prefix-list** `name route-targets` statement at the following hierarchy levels:

- [edit logical-systems *logical-system-name* policy-options]
- [edit logical-systems *logical-system-name* policy-options policy-statement *policy-name* term *term-name*]
- [edit policy-options]
- [edit policy-options policy-statement *policy-name* term *term-name*]

The following route target filtering match conditions are available:



**NOTE:** You define these match conditions in the **from** statement.

- **family route-target**—Specifies matching BGP route target filtering routes.
- **protocol route-target**—Specifies the criteria that an incoming route must match. This is useful for restricting the policy to locally generated route target filtering routes.
- **rtf-prefix-list name**—Applies the list of route target prefixes that you already configured to the policy.

As a result of the proxy BGP route target filtering feature, the **show route table bgp.rtarget.0** command has been updated to show the route target type of **Proxy**.

[[Understanding Proxy BGP Route Target Filtering](#)]

- **Static route target filtering**—Route target extended communities (see RFC 4360, *BGP Extended Communities Attribute*) prevent networks from receiving information about VPNs that is not relevant. For example, a network that does not include any PE routers that are a part of a VPN does not need to receive any network updates related to that VPN. The route target extended community feature has been extended to allow you to add static entries to the bgp.rtarget.0 routing table. This can be particularly useful for VPN hub-and-spoke topologies. Specify the target community for static route target

filtering using the **route-target-filter** statement at the **[edit routing-options rib bgp.rtarget.0 static]** hierarchy level.

[\[Configuring Static Route Target Filtering for VPNs\]](#)

**Related Documentation**

- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 154](#)
- [Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 170](#)
- [Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 171](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 285](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 322](#)

## Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers

### Changes in Default Behavior and Syntax

The following are changes made to Junos OS default behavior and syntax:

- [High Availability \(HA\) and Resiliency on page 154](#)
- [IPv6 on page 156](#)
- [Interfaces and Chassis on page 156](#)
- [J-Web on page 159](#)
- [Junos OS XML API and Scripting on page 159](#)
- [Layer 2 Ethernet Services on page 159](#)
- [MPLS on page 160](#)
- [Multicast on page 160](#)
- [Network Address Translation \(NAT\) on page 160](#)
- [Routing Protocols on page 161](#)
- [Security on page 162](#)
- [Services Applications on page 162](#)
- [Subscriber Access Management on page 162](#)
- [System Logging on page 169](#)
- [User Interface and Configuration on page 169](#)
- [VPNs on page 170](#)

#### ***High Availability (HA) and Resiliency***

- **Protection of MX, M, and T Series routers from denial of service (DOS) attacks**—New CLI options provide improved protection against DOS attacks.

- NAT mapping refresh behavior—Prior to this release, a conversation was kept alive when either inbound or outbound flows were active. This remains the default behavior. As of this release, you can also specify mapping refresh for only inbound flows or only outbound flows. To configure mapping refresh behavior, include the **mapping-refresh (inbound | outbound | inbound-outbound)** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.
- EIF inbound flow limit—Previously, the number of inbound connections on an EIF mapping was limited only by the maximum flows allowed on the system. You can now configure the number of inbound flows allowed for an EIF. To limit the number of inbound connections on an EIF mapping, include the **eif-flow-limit *number-of-flows*** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.

*[Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions]*

- **Determining readiness for graceful Routing Engine switchover in an MX Series Virtual Chassis (MX240, MX480, and MX960 routers with MPC/MIC interfaces)**—You can use the new **check** option for the **request virtual-chassis routing-engine master switch** command to determine whether the member routers in an MX Series Virtual Chassis configuration are ready for a global graceful Routing Engine switchover (GRES) operation from a database synchronization perspective. A global GRES changes the mastership in an MX Series Virtual Chassis by switching the global roles of the master router and backup router in the Virtual Chassis configuration.

Depending on the router configuration, a variable amount of time is required before a router is ready to perform a GRES operation. Attempting a GRES operation before the router is ready can cause system errors and unexpected behavior. Using the **request virtual-chassis routing-engine master switch check** command before you initiate the GRES operation ensures that the subscriber management and kernel databases on both member routers in an MX Series Virtual Chassis are synchronized and ready for the GRES operation.

The **request virtual-chassis routing-engine master switch check** command, which you must issue from the Virtual Chassis master router (VC-Mm), checks various system and database components to determine whether they are ready for GRES, but does not initiate the global GRES operation itself. The readiness check includes ensuring that a system timer, which expires after 300 seconds, has completed before the global GRES operation can begin.

If the member routers in an MX Series Virtual Chassis are ready for GRES from a database perspective, the **request virtual-chassis routing-engine master switch check** command returns the command prompt and displays no output. If the member routers are not ready for GRES, the command displays information about the readiness of the system.

*[Junos OS High Availability Configuration Guide]*

### IPv6

- **Change in automatically generated virtual-link-local-address for VRRP over IPv6—**Over IPv6 will be 0x02. This change makes the VRRP over IPv6 feature in Junos OS 12.2R5, 12.3R3, 13.1R3, and later releases inoperable with Junos OS 12.2R1, 12.2 R2, 12.2 R3, 12.2R4, 12.3R1, 12.3R2, 13.1R1, and 13.3R2 releases if an automatically generated virtual-link-local-address ID used. As a workaround, use a manually configured virtual-link-local-address instead of an automatically generated virtual-link-local-address.

### Interfaces and Chassis

- **New command to monitor PPP recovery after a GRES or restart (MX Series routers)—**The new **show ppp statistics recovery** command monitors the progress of PPP recovery after a GRES or restart. When the PPP subscriber sessions have been recovered, the command output displays **Recovery state: recovery done** to indicate that it is safe to force another GRES or restart. When you issue this command during the recovery process, the command might time out or fail silently rather than display output. Recovery is not complete until the command displays recovery done.

*[Interfaces Command Reference]*

- **Enhancement to set date ntp command—**You can now specify an authentication-key number for the NTP server used to synchronize the date and time on the router or switch. Include the new **key number** option with the **set date ntp** command. The key number you include must match the number you configure for the NTP server at the **[edit system ntp authentication-key number]** hierarchy level.

*[System Basics and Services Command Reference]*

- **New fast-failover option for LACP—**You can now configure the Link Aggregation Control Protocol for aggregated Ethernet interfaces to facilitate subsecond failover. To override the default behavior for the IEEE 802.3ad standard and allow the standby link always to receive traffic, include the **fast-failover** statement at the **[edit interfaces aex aggregated-ether-options lacp]** hierarchy level.

*[Junos OS Ethernet Interfaces Configuration Guide]*

- **Inclusion of AC-Name and AC-Cookie tags in PPPoE PADS messages by default—**By default, a router that functions as an access concentrator (AC) sends both the AC-Name and AC-Cookie tags as part of the PPPoE Active Discovery Session (PADS) packet when it confirms a session with a PPPoE client. You can configure the **no-send-pads-ac-info** statement at the **[edit protocols pppoe]** hierarchy level to prevent the router from transmitting the AC-Name and AC-Cookie tags in the PADS messages.

*[Network Interfaces, Ethernet Interfaces]*

- **Enhancement to Link Layer Discovery Protocol (LLDP) (MX Series and T Series routers)—**You can configure LLDP to generate the interface name as the port ID Type TLV. To generate the interface name as the port ID Type, Length, and Value, include the **interface-name** statement at the **[edit protocols lldp port-id-subtype]** hierarchy level. The default behavior is to generate the SNMP Index of the interface as the port ID TLV. If you have changed the default behavior, include the **locally-assigned** statement

at the **[edit protocols lldp port-id-subtype]** hierarchy level to reenab the default behavior of generating the SNMP Index of the interface as the port ID TLV. When you configure LLDP to generate the interface name as the port ID TLV on the remote neighbor, the **show lldp neighbors** command displays the interface name in the **Port ID** field. The default behavior is for the command to display the SNMP index of the interface of the remote neighbor in the **Port ID** field. [*Ethernet Interfaces Configuration Guide, Interfaces Command Reference*]

- **New Link Aggregation Control Protocol (LACP) Commands and SNMP MIB**--You can now view and clear LACP timeout entries. To display information about LACP timeout entries, use the **show lacp timeouts** command. Include the **interfaces interface-name** option to view timeout information about a specific interface only. To clear LACP timeout entries, use the **clear lacp timeouts** command. Include the **interfaces interface-name** option to clear timeout information for a specific interface only. A new SNMP MIB is now also available. The **jnxLacpAggTimeout** MIB lists all interfaces where the **jnxLacpTimeOut** trap is sent. [*Interfaces Command Reference*]
- **Enhancement to show interfaces queue command**—The output for the **show interfaces queue** command now displays rate-limit statistics for class-of-service schedulers for all IQ and Enhanced IQ (IQ2E) PICs when rate-limiting is configured, even when no traffic is dropped. When rate limiting is configured but no traffic is dropped, the output for the **RL-dropped packets** and **RL-dropped-bytes** fields display the value zero (0). Previously, these fields were not displayed when no traffic was dropped and rate-limiting was configured. To configure rate-limiting for queues before packets are queued for output, you include the **rate-limit** statement at the **[edit class-of-service schedulers transmit-rate rate]** hierarchy level. [*Interfaces Command Reference*]
- On MX80 routers, the FPC Slot output field has been changed to TFEB Slot for the **show services accounting flow inline-jflow**, **show services accounting errors inline-jflow**, and **show services accounting status inline-jflow** commands.
- **Configuring the flow-tap service for IPv6 traffic**—The **family inet | inet6** statement at the **[edit services flow-tap]** hierarchy level enables you to specify the type of traffic for which you want to apply the flow-tap service. If the family statement is not included, the flow-tap service is, by default, applied to the IPv4 traffic. To apply flow-tap service to IPv6 traffic, you must include the **family inet6** statement in the configuration. To enable the flow-tap service for IPv4 and IPv6 traffic, you must explicitly configure the family statement for both inet and inet6 families.

However, you cannot configure the flow-tap service for IPv6 along with port mirroring or sampling of IPv6 traffic on routers that support LMNR-based FPCs. This restriction is valid even if the router does not have an LMNR-based FPC installed on it. There is no restriction on configuring the flow-tap service on routers that are configured for port mirroring or sampling of IPv4 traffic.

- Prior to Junos OS Release 12.2, when you issue the **show system memory** command on MX80 routers, the **unable to load pmap\_helper module: No such file or directory** error message is displayed in the output of the command. Starting with Junos OS Release 12.2, PMAP information is correctly displayed in the output of this command for MX80 and ACX Series routers.

[*System Basics and Services Command Reference*]

- **Changes to DDoS protocol groups (MX Series routers)**—The **ipv4-unclassified** and **ipv6-unclassified** DDoS protocol groups have been deprecated in the **protocols** statement at the **[edit system ddos-protection ddos]** hierarchy level. These two protocol groups have also been deprecated from the **show ddos-protection protocols** commands. These groups formerly were used to police all unclassified IPv4 and IPv6 host-bound traffic.

In their place, 10 new protocol groups have been added to the **protocols** statement and the **show ddos-protection protocols** commands:

- **control-layer2**—Unclassified layer 2 control packets.
- **control-v4**—Unclassified IPv4 control packets.
- **control-v6**—Unclassified IPv6 control packets.
- **filter-v4**—Unclassified IPv4 filter action packets; sent to the host because of reject terms in firewall filters.
- **filter-v6**—Unclassified IPv6 filter action packets; sent to the host because of reject terms in firewall filters.
- **host-route-v4**—Unclassified IPv4 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
- **host-route-v6**—Unclassified IPv6 routing protocol and host packets in traffic sent to the router local interface address for broadcast and multicast.
- **other**—All unclassified packets that do not belong to another type.
- **resolve-v4**—Unclassified IPv4 resolve packets sent to the host because of a traffic request resolve action.
- **resolve-v6**—Unclassified IPv6 resolve packets sent to the host because of a traffic request resolve action.

*[DDoS Configuration]*

- **New range for message-rate-limit**—The range for message-rate-limit under the syslog configuration for services has changed to 0 through 2147483647.
- **SNMP Traps for FPC crash events (T Series)**—The **jnxFruTable** object (in the Chassis MIB) is supported for FPC crash events on T Series routers. You can use the **show log messages | match trap** command to view the SNMP Traps.
- **SNMP Traps for SPMB crash events (T Series)**—The **jnxFruTable** object (in the Chassis MIB) is supported for SPMB (Switch Processor Mezzanine Board) crash events on T Series routers. You can use the **show log chassisd** command to view the SNMP MIB objects.

### *J-Web*

- On all M Series, MX Series, and T Series platforms, the username field does not accept HTML tags or the < and > characters. The following error message appears: **A username cannot include certain characters, including < and >.**

### *Junos OS XML API and Scripting*

- **IPv6 address text representation is stored internally and displayed in command output using lowercase**—Starting with Junos OS Release 11.1R1, IPv6 addresses are stored internally and displayed in the command output using lowercase. Scripts that match on an uppercase text representation of IPv6 addresses should be adjusted to either match on lowercase or perform case-insensitive matches.
- **<get-configuration> RPC with inherit="inherit" attribute returns correct time attributes for committed configuration**—In Junos OS Release 12.2R3 and earlier releases, when you configured some interfaces using the interface-range configuration statement, if you later requested the committed configuration using the <get-configuration> RPC with the inherit="inherit" and database="committed" attributes, the device returned junos:changed-localtime and junos:changed-seconds in the RPC reply instead of junos:commit-localtime and junos:commit-seconds. This issue is fixed in Junos OS Release 12.2R4 and later releases so that the device returns the expected attributes in the RPC reply.

### *Layer 2 Ethernet Services*

- **Support for displaying logical system and routing instance for L2TP tunnels (MX Series routers)**—When you issue the **show services l2tp tunnel** command with the **detail** or **extensive** option on either the LAC or LNS, the output now displays both the logical system and the routing instance in which the L2TP tunnel is brought up.

*[System Basics and Services Command Reference]*

- **New options for Multichassis Link Aggregation (MC-LAG)**—For MC-LAG, you can now specify one of two actions to take if the Inter-Chassis Communication Protocol (ICCP) peer if the switch or router goes down. To bring down the interchassis link logical interface if the peer goes down, include the **force-icl-down** statement at the **[edit interfaces aeX aggregated-ether-options events iccp-peer-down]** hierarchy level. To have the router or switch become the active node when a peer goes down, include the **prefer-status-control-active** statement at the **[edit interfaces aeX aggregated-ether-options mc-ae events iccp-peer-down]** hierarchy level. When you configure the **prefer-status-control-active** statement, you must also configure the **status-control active** statement at the **[edit interfaces aeX aggregated-ether-options-mc-ae]** hierarchy level. If you do not configure the **status-control** as **active** with the **prefer-status-control-active** statement, the router or switch does not become the active node if a peer goes down. *[Junos OS Ethernet Interfaces Configuration Guide]*
- **Multichassis Link Aggregation (MC-LAG)**—When you configure the **prefer-status-control-active** statement at the **[edit interfaces aex aggregated-ether-options mc-ae events iccp-peer-down]** hierarchy level, you must also configure the **status-control active** statement at the **[edit interfaces aex**

**aggregated-ether-options mc-ae**] hierarchy level. If you configure the **status-control standby** statement with the **prefer-status-control-active** statement, the system issues a warning. [*Junos OS Ethernet Interfaces Configuration Guide*]

### **MPLS**

- **Policers for MPLS LSPs (T Series Core Routers)**—You can now configure an MPLS LSP policer for a specific LSP to be shared across different protocol family types. To do so, you must configure the LSP policer as a logical interface policer. Include the **logical-interface-policer** statement at the **[edit firewall policer policer-name]** hierarchy level. Previously, you could not configure an MPLS LSP policer as a logical interface policer. When you configure an MPLS LSP policer as a logical interface policer, that single policer polices traffic for all protocol families for an LSP. An MPLS LSP policer not configured as a logical interface policer continues to police traffic for a specific protocol family only. [*Firewall Filters and Traffic Policers Configuration Guide, MPLS Applications Configuration Guide*]
- Starting in Junos OS Release 12.2, at the end of each **adjust-interval**, LSP's **max\_average** for the **auto-bandwidth** functionally does not reset to zero. The **max\_average** retains the value from the last interval until the first sample of the current interval is received. When the first sample of the current interval is received, the **max\_average** is updated to the first sample value.

In the **show mpls lsp** command output, the value for **Max AvgBW util** now displays the value of the maximum average bandwidth utilization from the previous interval until the first sample of the current interval is obtained.

[*MPLS Operational Mode Commands*]

### **Multicast**

- Starting in Junos OS Release 8.0, the TTL value for PIM Graft messages, which are unicast, is set to 1. Previously, the TTL for PIM Graft messages was set to 64.
- In a bootstrap router (BSR)-enabled bidirectional PIM domain, mixing Junos OS Release pre-12.1R7 releases and later releases can cause unexpected outages. If you have a deployment with routers running Junos OS Release pre-12.1R7 and if you upgrade a subset of the routers to Junos OS Release 12.1R7 or later, the group-to-RP mapping across the domain breaks and an outage occurs.

### **Network Address Translation (NAT)**

- **Limitation on number of terms for NAT rules applied to inline services interfaces**—You are limited to a maximum of 200 for a NAT rule that is applied to an inline services (type **si**) interface. If you specify more than 200 terms, you will receive the following error when you commit the configuration:

```
[edit]
'service-set service-set-name'
  NAT rule rule-name with more than 200 terms is disallowed for si-x/y/z.n
error: configuration check-out failed
```



- The method for computing the block size for deterministic port block allocation for network port translation (NAPT) when the configured block size is zero has changed, and is computed as follows:

$$\text{block-size} = \text{int}(64512 / \text{ceil}[(\text{Nr\_Addr\_PR\_Prefix} / \text{Nr\_Addr\_PU\_Prefix})])$$

where:

64512 is the maximum available port range per public IP address.

Nr\_Addr\_PR\_Prefix is the number of usable pre-NAT IPv4 subscriber addresses in a **from** clause match condition.

Nr\_Addr\_PU\_Prefix is the number of usable post-NAT IPv4 addresses configured in the NAT pool.

### ***Routing Protocols***

- **Support for processing large PDUs in IS-IS:**

- New option to disable hello padding on IS-IS packets—The **hello-padding** statement has a new option, **disable**, which can be used to disable padding of hello packets on all types of interfaces for all adjacency states.
- New statement to limit size of IS-IS hello packets—The **max-hello-size size** statement is introduced at the **[edit protocols isis]** hierarchy level to modify the maximum size of IS-IS hello packets. The size varies from 512 through 1492 bytes. The default size is 1492 bytes.
- New statement to limit the size of IS-IS link-state PDUs—The **max-lsp-size size** statement is introduced at the **[edit protocols isis]** hierarchy level to modify the maximum size of IS-IS link-state PDUs. The size varies from 512 through 1492 bytes. The default size is 1492 bytes.
- New statement to limit the size of IS-IS sequence number packets—The **max-snp-size size** statement is introduced at the **[edit protocols isis]** hierarchy level to modify the maximum size of partial or complete IS-IS sequence number packets. The size varies from 512 through 1400 bytes. The default size is 1400 bytes.

*[Routing Protocols Configuration Guide]*

- This release supports a new **show firewall templates-in-use** operational command. This command enables you to display the names of filters configured using the filter statement at either the **[edit firewall]** or **[edit dynamic-profiles profile-name firewall]** hierarchy level and that are being used as templates for dynamic subscriber filtering. The command also displays the number of times the filter has been referenced by subscribers accessing the network.

*[Routing Protocols and Policies Command Reference]*

- If you configure the **route-distinguisher** statement in addition to the **route-distinguisher-id** statement, the value configured for **route-distinguisher** supersedes the value generated from **route-distinguisher-id**. To avoid a conflict in the two route distinguisher values, we recommend that you ensure the first half of the route distinguisher obtained by configuring the **route-distinguisher** statement is different from the first half of the route distinguisher obtained by configuring the **route-distinguisher-id** statement.

- When configuring the **advertise-external** statement for an AS confederation, we recommend that EBGP peers belonging to different autonomous systems are configured in a separate EBGP peer group. This ensures consistency while BGP sends the best external route to peers in the configured peer group.
- Prior to Junos OS Release 12.2, groups of peer bits in the output of the **show bgp group rtf detail** command were displayed in reverse order.
- Starting in Junos OS Release 12.2, the **show bgp group** output is updated to a new multiline format in order to display the full name of table `bgp.rtarget.0`.

### Security

- In all supported Junos OS releases, regular expressions can no longer be configured if they require more than 64MB of memory or more than 256 recursions for parsing.

This change in the behavior of Junos OS is in line with the Free BSD limit. The change was made in response to a known consumption vulnerability that allows an attacker to cause a denial of service (resource exhaustion) attack by using regular expressions containing adjacent repetition operators or adjacent bounded repetitions. Junos OS uses regular expressions in several places within the CLI. Exploitation of this vulnerability can cause the Routing Engine to crash, leading to a partial denial of service. Repeated exploitation can result in an extended partial outage of services provided by the routing protocol process (`rpd`).

### Services Applications

- Starting in Junos OS Release 12.2R5, the **destination-address** statement in a firewall rule **from** statement might not have the address value of `0::00` with IPv6.

```
[edit services stateful-firewall rule rule-name term term-name from]
destination-address (address | any-unicast) <except>;
```

This issue is being tracked by [PR857106](#).

- **Restrictions for maximum blocksize for NAT port block allocation**—The maximum blocksize for NAT port block allocation (PBA) is now 32,000.

### Subscriber Access Management



**NOTE:** Although present in the code, the subscriber management features are not supported in Junos OS Release 12.2R9. Documentation for subscriber management features is included in the Junos OS Release 12.2 documentation set.

---

- **Additional option for RADIUS NAS-Port attribute (MX Series routers)**—You can now configure the width of the aggregated Ethernet identifier field used in the RADIUS NAS-Port attribute (attribute 5). To configure the width, include the **ae-width** option in the **nas-port-extended-format** statement at the **[edit access profile *profile-name* radius options]** hierarchy level. The **ae-width** field can be from 0 through 32 bits. The total width of the NAS-Port attribute can be a maximum of 32 bits.

[Subscriber Access]

- **Making Ascend-Data-Filter optional for dynamic subscribers (MX Series routers)**—When you configure the \$junos-adf-rule-v4 or \$junos-adf-rule-v6 variable for an Ascend-Data-Filter in a dynamic profile, an error is reported when the RADIUS reply message does not include the variable for subscriber sessions affected by the dynamic profile. Consequently, system resource utilization is increased when the dynamic profile is applied to a mix of subscribers where RADIUS does not associate some of the subscribers with an Ascend-Data-Filter.

In this situation, you can reduce the effect on system resources by including the **not-mandatory** option in the Ascend-Data-Filter configuration at the **[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-unit-number* family *family* filter adf]** hierarchy level. The **not-mandatory** option suppresses error reporting when the variable is not present in the RADIUS message and prevents the Ascend-Data-Filter from being created.

- **Support for controlling the negotiation order of PPP authentication protocols (MX Series routers)**—You can control the order in which the router tries to negotiate PPP authentication protocols when it verifies that a PPP client can access the network. By default, the router tries to negotiate Challenge Handshake Authentication Protocol (CHAP) authentication first, and then tries Password Authentication Protocol (PAP) if the attempt to negotiate CHAP authentication is unsuccessful. You can now modify the default negotiation order for CHAP and PAP to suit your subscriber network requirements.

In earlier Junos OS releases, you could not change the default negotiation order for CHAP and PAP. The router always tried negotiating CHAP authentication first, followed by PAP authentication if CHAP negotiation was unsuccessful.

To configure the negotiation order for CHAP and PAP authentication, issue the new **authentication** statement at the **[edit dynamic-profiles *profile-name* interfaces *pp0* unit “\$junos-interface-unit” ppp-options]** hierarchy level (for dynamic PPP subscriber interfaces) or at the **[edit interfaces *pp0* unit *logical-unit-number* ppp-options]** hierarchy level (for static interfaces with PPP encapsulation).

You can issue the **authentication** statement in any of the following ways:

- To specify that the router negotiate PAP authentication first, followed by CHAP authentication if PAP negotiation is unsuccessful, issue the **authentication [pap chap]** statement. When you specify both authentication protocols in either order, you must enclose the set of protocol names in square brackets ([ ]).
- To specify that the router negotiate only CHAP authentication, issue the **authentication chap** statement.

- To specify that the router negotiate only PAP authentication, issue the **authentication pap** statement.

[Subscriber Access]

- **Support for modifying the CHAP challenge length (MX Series routers)**—You can modify the default minimum length and maximum length of the Challenge Handshake Authentication Protocol (CHAP) challenge message that the router sends to a PPP client. By default, the minimum length of the CHAP challenge is 16 bytes, and the maximum length is 32 bytes. You can override this default to configure the CHAP challenge minimum length and maximum length in the range 8 bytes through 63 bytes.



**BEST PRACTICE:** We recommend that you configure both the minimum length and the maximum length of the CHAP challenge to at least 16 bytes.

In earlier Junos OS releases, you could not change the default length of the CHAP challenge message.

To configure the minimum and maximum length of the CHAP challenge message, issue the new **challenge-length** statement at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit “\$junos-interface-unit” ppp-options chap]** hierarchy level (for dynamic PPP subscriber interfaces) or at the **[edit interfaces pp0 unit *logical-unit-number* ppp-options chap]** hierarchy level (for static interfaces with PPP encapsulation).

For example, the following **challenge-length** statement in a dynamic profile named **pppoe-client-profile** sets the minimum length of the CHAP challenge to 20 bytes, and the maximum length to 40 bytes.

```
[edit dynamic-profiles pppoe-client-profile interfaces pp0 unit “$junos-interface-unit”  
  ppp-options chap]  
user@host# set challenge-length minimum 20 maximum 40
```

[Subscriber Access]

- **Support for agent circuit identifier filtering in PPPoE subscriber session lockout (M120, M320, and MX Series routers)**—Extends the PPPoE subscriber session lockout feature, which is also referred to as PPPoE encapsulation type lockout, to support identification and filtering of PPPoE subscriber sessions by either the agent circuit identifier (ACI) value or the unique media access control (MAC) source address on the PPPoE underlying interface. In earlier Junos OS releases, you used PPPoE subscriber session lockout to identify and filter subscriber sessions only by their unique MAC source address.

Configuring and using PPPoE subscriber session lockout increases router efficiency and protects the router and any external AAA servers from excessive loading by temporarily deferring failed or short-lived subscriber sessions in favor of those sessions that can complete successfully. PPPoE subscriber session lockout enables you to prevent (lock out) a failed or short-lived PPPoE subscriber session from reconnecting to the router for a default or configurable period of time, based on either of the following options:

- The subscriber session's unique MAC source address on the PPPoE underlying interface

This option, which is the default, locks out the offending PPPoE subscriber session identified by its MAC source address on the underlying interface.

- The ACI string contained in the DSL Forum Agent-Circuit-ID VSA [26-1] (option 0x105) of PPPoE Active Discovery Initiation (PADI) and PPPoE Active Discovery Request (PADR) control packets

This option locks out all PPPoE subscriber sessions on the underlying interface that come from the same household and share the same ACI string in their PPPoE PADI and PADR control packets.

PPPoE subscriber session lockout based on the ACI value is particularly useful for configurations such as the following in which MAC source addresses are not unique on the PPPoE underlying interface:

- PPPoE interworking function sessions in which the MAC addresses of all PPPoE interworking function sessions contain the MAC address of the DSLAM device
- Configurations in which the access node (usually a DSLAM device) overwrites the MAC source address in PPPoE packets received from the customer premises equipment (CPE) with its own MAC address for security purposes
- Duplicate MAC source addresses across disparate households in an N:1 (service VLAN) configuration, which requires the router to use a combination of the MAC source address and the ACI value to uniquely identify a subscriber

To configure temporary PPPoE subscriber session lockout based on the ACI value, include the **short-cycle-protection** statement with the new **filter aci** option for PPPoE subscriber sessions on any of the following underlying logical interfaces types:

- Dynamic or static VLAN interfaces (in the **pppoe-underlying-options** stanza)
- Dynamic or static VLAN demultiplexing (demux) interfaces (in the **family pppoe** stanza)

For example, the following statement configures temporary lockout based on ACI information for PPPoE subscriber sessions on a dynamic VLAN underlying interface. This statement specifies a nondefault lockout time in the range 20 through 120 seconds.

```
[edit dynamic-profiles my-vlan-profile interfaces "$junos-interface-ifd-name" unit  
"$junos-interface-unit" pppoe-underlying options]  
user@host# set short-cycle-protection lockout-time-min 20 lockout-time-max 120  
filter aci
```

The following statement configures temporary lockout based on ACI information for PPPoE subscriber sessions on a dynamic VLAN demux underlying interface. This statement uses the default lockout time range 1 through 300 seconds.

```
[edit dynamic-profiles my-demux-vlan-profile interfaces demux0 unit
 "$junos-interface-unit" family pppoe]
user@host# set short-cycle-protection filter aci
```

The **clear pppoe lockout** operational command has been enhanced in this release to clear the lockout condition for PPPoE subscriber sessions associated with a particular ACI value. For example, the following command clears the lockout condition for all PPPoE subscriber sessions on underlying VLAN demux interface demux0.214 associated with an ACI value that matches the regular expression "Relay-identifier atm 3/0:100.\*". You must enclose the regular expression in quotation marks.

```
user@host> clear pppoe lockout underlying-interfaces demux0.214 aci "Relay-identifier
 atm 3/0:100.*"
```

To display information about PPPoE subscriber session lockout based on ACI information, you can also use the enhanced **show pppoe lockout** and **show pppoe underlying-interfaces** operational commands.

[Subscriber Access]

- The *Example: HTTP Service Attached to a Static Interface* topic in the *Junos OS Subscriber Access Configuration Guide* provides an incorrect example for configuring a service filter as a walled garden. The correct example is as follows:

The following example uses a service filter as a walled garden by defining a rule named **redirect**, referencing the rule in a profile named **http-redirect**, configuring a service set named **http-redirect** that references the **http-redirect** captive portal content delivery profile, and attaching the **http-redirect** service set to static interface **ge-1/0/1.0**.

```
[edit services]
captive-portal-content-delivery {
  rule redirect {
    match-direction input;
    term t1 {
      from {
        destination-address {
          100.0.1.1/32;
        }
      }
      then {
        redirect http://www.google.com;
      }
    }
  }
  profile http-redirect {
    cpcd-rules redirect;
  }
}
service-set http-redirect {
  captive-portal-content-delivery-profile http-redirect;
  interface-service {
    service-interface ms-1/0/0;
```

```

    }
  }
  [edit interfaces ge-1/0/1]
  unit 0 {
    family inet {
      service {
        input {
          service-set http-redirect service-filter walled;
        }
        output {
          service-set http-redirect;
        }
      }
    }
    address 10.1.3.2/24;
  }
}

```

[Subscriber Access]

- On MX80 routers, you can configure only four inline services physical interfaces as anchor interfaces for L2TP LNS sessions: si-1/0/0, si-1/1/0, si-1/2/0, si-1/3/0. You cannot configure si-0/0/0 for this purpose on MX80 routers.
- The **user *username*** option for the **clear services l2tp session** command is no longer available in the CLI for LNS on MX Series routers. Added to the option's previous unavailability for LAC on MX Series routers, this means that L2TP on MX Series routers does not support clearing L2TP sessions based on subscriber username. As an alternative, you can determine the session ID for the username by issuing the **show subscribers detail** command, and then remove the session with the **clear services l2tp session local-session-id *session-id*** command.

[Subscriber Access]

- The **user *username*** option for the **show services l2tp session** command is no longer available in the CLI for L2TP LAC or L2TP LNS on MX Series routers. To view L2TP session information organized by subscriber username, you can issue the **show subscribers detail** command or the **show network-access aaa subscribers username** command.

[Subscriber Access]

- **Effect of changing the forwarding class configuration with PPP fast keepalive (MX Series routers with MPC/MIC interfaces)**—To change the default queue assignment (forwarding class) for outbound traffic generated by the Routing Engine, you can include the **forwarding-class *class-name*** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level.

For PPP fast (inline) keepalive LCP Echo-Request and LCP Echo-Reply packets transmitted between an MX Series router with MPCs/MICs and a PPP client, changing the forwarding class configuration takes effect immediately for both new PPP-over-Ethernet (PPPoE), PPP-over-ATM (PPPoA), and L2TP network server (LNS) subscriber sessions created after the configuration change, and for existing PPPoE, PPPoA, and LNS subscriber sessions established before the configuration change.

In earlier Junos OS releases with PPP fast keepalive, forwarding class configuration changes applied only to new PPPoE, PPPoA, and LNS subscriber sessions created after the configuration change. The forwarding class setting was fixed for existing PPPoE, PPPoA, and LNS subscriber sessions, and could not be changed until the session was terminated and re-established.

[*Junos OS Subscriber Access Configuration Guide, Junos OS Class of Service Configuration Guide*]

- When an MX Series router configured as an LNS sends an Access-Request message to RADIUS for an LNS subscriber, the LNS now includes the Called-Station-ID-Attribute when it receives AVP 21 in the ICRQ message from the LAC.
- **L2TP support for SNMP statistics (MX Series routers)**—By default, SNMP polling is disabled for L2TP statistics. As a consequence, the L2TP tunnel and global counters listed in the table have a default value of zero.

**Table 3: SNMP Counters for L2TP Statistics**

Counter Name	Type
jnxL2tpTunnelStatsDataTxPkts	Tunnel
jnxL2tpTunnelStatsDataRxPkts	Tunnel
jnxL2tpTunnelStatsDataTxBytes	Tunnel
jnxL2tpTunnelStatsDataRxBytes	Tunnel
jnxL2tpStatsPayloadRxOctets	Global
jnxL2tpStatsPayloadRxPkts	Global
jnxL2tpStatsPayloadTxOctets	Global
jnxL2tpStatsPayloadTxPkts	Global

You can enable collection of these statistics by including the **enable-snmp-tunnel-statistics** statement at the **[edit services l2tp]** hierarchy level. When enabled, the L2TP process polls for these statistics every 30 seconds for 1000 sessions. The potential age of the statistics increases with the number of subscriber sessions; the data is refreshed more quickly as the number of sessions decreases. For example, with 30,000 sessions, none of these statistics is more than 15 minutes old.



**BEST PRACTICE:** The system load can increase when you enable these counters and also use RADIUS interim accounting updates. We recommend you enable these counters when you are using only SNMP statistics.

[*Subscriber Access*]



- **Updated AAA Terminate Reason Mappings (MX Series routers)**—The AAA **idle-timeout** terminate reason is now mapped to the RADIUS accounting Idle Timeout (4) terminate cause, and the AAA **session-timeout** terminate reason is now mapped to the RADIUS Session Timeout (5) terminate cause. In earlier releases, both terminate reasons were mapped to the RADIUS accounting NAS Request (10) terminate cause.

To support backward compatibility, you can configure the router to support the previous behavior—use the **terminate-code aaa shutdown (idle-timeout | session-timeout) radius 10** statement at the **[edit access]** hierarchy level.

[Subscriber Access]

- **DHCP client IP address (MX Series)**—Starting in Junos OS Release 12.2, you can configure the subnet to which the DHCP local server matches the requested IP address. The server accepts and uses an active client's requested IP address to address assignment only when the requested address and the IP address of the DHCP server interface are in the same subnet. The server accepts and uses a passive client's requested IP address only when the requested address and the IP address of the relay interface are in the same subnet.

### System Logging

- Prior to Junos OS Release 12.2, when a downstream non-Juniper Networks router sent an incorrect RESV message with a bandwidth in FlowSpec that was less than the bandwidth required by TSpec in the Path message, the following warning message was logged - **RPD\_RSVP\_INCORRECT\_FLOWSPEC**. Starting with Junos OS Release 12.2, the **RPD\_RSVP\_INCORRECT\_FLOWSPEC** error message is not logged, as a peak rate mismatch does not affect router functionality.

[System Log]

- **Enhancements to Cannot perform nh operation ADDANDGET system log message**—Prior to Junos OS Release 11.4, the **Cannot perform nh operation ADDANDGET** system log message was getting logged many times while bringing up clients in an MX Series subscriber services deployment, which impacted system performance. To prevent the **Cannot perform nh operation ADDANDGET** system log message from being logged multiple times, starting with Junos OS Release 11.4, the message is rate limited. Besides rate-limiting the message, additional information, such as, **nhindex**, **ifindex**, **fwdnhidx**, and the number of suppressed logs is also displayed in the log message.

The following is a sample of the enhanced system log message:

```
Jun 13 14:00:00 calcium rpd[1332]: Cannot perform nh operation ADDANDGET nhop
0.0.0.0 type unicast nhindex 0x0 ifindex 0x1471 demux0.1073991785 fwdnhidx 0x0
type unicast errno 45 suppressed 412 logs
```

[System Log]

### User Interface and Configuration

- **Enhancement to test configuration operational statement**— The option **syntax-only** allows a user to check a partial configuration without checking for commit errors.

[System Basics and Services Command Reference]

- **Removal of the sampling action modifier for IPv4 firewall filters**—In the J-Web interface, the **Sample** check box is not available for configuration from the Other Actions section under the Actions tab of the Configure > Security > Filters > IPv4 Firewall Filters page. This configuration of the sample action modifier is not enabled because the **set firewall filter foo term bar then sample** configuration command has been deprecated in the Junos OS CLI in Release 12.1 and later.

*[J-Web Online Help]*

#### VPNs

- L2VPN/BGP-VPLS site-preference's value has changed its encoding formatting. Value has been traditionally encoded with an erroneous endiannes format, which prevented Junos OS platforms from inter-operating with other vendors. While the current fix corrects this behavior, it should be taken into account, that updated Junos OS platforms (containing this fix) may be unable to interoperate with those ones that are running older releases.

#### Related Documentation

- [New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 100](#)
- [Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 170](#)
- [Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 171](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 285](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 322](#)

### Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 12.2R9 for the M Series, MX Series, and T Series.

For the most complete and latest information about known Junos OS defects, use the Juniper Networks online [Junos Problem Report Search](#) application.

- [Class of Service \(CoS\)](#)
- [Subscriber Management and Services](#)

#### Class of Service (CoS)

- If you define more than one forwarding class for a given queue number, do not use the name of a default forwarding class for one of the new classes, because doing so causes the forwarding class with the default name to be deleted. For example, do not configure the following, because doing so deletes the **best-effort** class:

```
user@host# set class-of-service forwarding-classes class be queue-num 0
user@host# set class-of-service forwarding-classes class best-effort queue-num 0
user@host# commit
```

### Subscriber Management and Services

---

- On MX Series, subscriber management uses firewall filters to capture and report the volume-based service accounting counters that are used for subscriber billing. You must always consider the relationship between firewall filters and service accounting counters, especially when clearing firewall statistics. When you use the **clear firewall** command (to clear the statistics displayed by the **show firewall** command), the command also clears the service accounting counters that are reported to the RADIUS accounting server. For this reason, you must be cautious in specifying which firewall statistics you want to clear. When you reset firewall statistics to zero, you also zero the counters reported to RADIUS.

#### Related Documentation

- [New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 100](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 154](#)
- [Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 171](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 285](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 322](#)

## Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers

The current software release is Release 12.2. For information about obtaining the software packages, see “[Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers](#)” on page 322.



**CAUTION:** This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for M, MX, and T Series routers. If you use unified ISSU when upgrading between releases, see the outstanding issues for [High Availability \(HA\) and Resiliency on page 214](#) in the current software release.

- 
- [Current Software Release on page 172](#)
  - [Previous Releases on page 201](#)

## Current Software Release

---

### Outstanding Issues

#### Class of Service (CoS)

- M7i and M10i without enhanced CFEB only supports four forwarding classes. However, when the user configured more than four forwarding classes, commit check didn't throw any errors, and additionally "show interfaces queue" showing supported FCs as 16. [PR786081](#)
- When FPC/PIC restarts or Routing Engine reboots, the physical interfaces are created and Class of Service daemon (cosd) sends chassis scheduler ADD for all interfaces. If a group of physical interfaces share the same Packet Forwarding Engine stream (such as oversubscribed PIC PD-5-10XGE-SFPP) and user configured chassis-scheduler is applied on some (NOT all) of the interfaces, the user configured chassis-scheduler can get over-written by default scheduler when chassis scheduler ADD comes for other non configured interface in same stream group with default scheduler-map. The issue can happen on any queuing PIC where multiple physical interfaces on PIC share same Packet Forwarding Engine/Chassis stream on FPC, in this bug, the fix is ONLY for PD-5-10XGE-SFPP. [PR809528](#)
- The names "best-effort", "assured-forwarding", "expedited-forwarding", "network-control" are reserved and cannot be currently used in Forwarding Class alias configuration, with several classes mapped to the same queue: `user@router# show class-of-service user@router# set class-of-service forwarding-classes class best-effort queue-num 0 user@router# set class-of-service classifiers inet-precedence test forwarding-class best-effort loss-priority low code-points 000 user@router# commit check configuration check succeeds user@router# set class-of-service forwarding-classes class myBE queue-num 0 user@router# commit check [edit class-of-service classifiers inet-precedence test forwarding-class] 'best-effort' forwarding class undefined: best-effort error: configuration check-out failed. PR827496`
- COSD errors are seen while Routing Engine switchover without GRES enabled. [PR827534](#)
- COSD errors - COSD\_GENCFG\_WRITE\_FAILED: GENCFG write failed (op, minor\_type) = (add, policy inline) for tbl 4 if 7454 /2/0 Reason: File exists are during Routing Engine switchover. [PR827538](#)
- In PPPoE/DHCP subscriber management environment, with "burst-size \$junos-cos-shaping-rate-burst" configured in subscriber dynamic-profiles, while logging in/out subscribers, the class-of-service daemon (cosd) memory leak due to cosd process doesn't free up memory used for parsing burst attributes of a traffic-control-profiles (tcp) guaranteed rate. The memory usage of cosd process can be monitored by following CLI command: `user@router> show system processes extensive | match "PID | cosd"` (Note: The "RES" field means "Current amount of resident memory, in kilobytes")  
PID USERNAME THR PRI NICE SIZE RES STATE TIME  
WCPU COMMAND 1326 root 1 96 0 14732K 4764K select 0:01 0.00% cosd. [PR846615](#)
- The output of the show subscribers extensive command displays the Effective shaping-rate field only if you have enabled the effective shaping rate at the [edit chassis] hierarchy level. [PR936253](#)

- Applying a scheduler with transmit rate below 65,535 bps and rate-limit option fails the commit if the associated interface is a non-existing interface or a virtual interface. [PR964647](#)
- CoS relevant misconfiguration (e.g. configure classifier exp for LT interfaces implicitly using "interface all" way) might cause cosd crash. If cosd experiences multiple crashes within a short time, it might not be able to restart. [PR969900](#)

### ***Forwarding and Sampling***

- Firewall logs do not record rejected packets on the loopback filter though packets are rejected. [PR724059](#)
- "show firewall detail" command will not display all the firewall policer counters if the "enhanced-policer" chassis knob is set and if the configured filters contain more than 10 policer counters. [PR789889](#)
- This is a cosmetic issue. If we prepare following conditions, we can find this behavior when we delete interface policer configuration. We cannot see this behavior without "commit synchronize". < Conditions > 1. Use 64bit Junos. 2. Configure "graceful-switchover" and "policer". 3. Delete interface policer configuration and then hit "commit synchronize". < backup RE messages > dfw\_update\_local\_shared\_policer: new filter program should be NULL for op 3 If you find this issue with fixed code, please re-configure "system syslog". [PR873084](#)
- In T4000 platforms with ES-FPC, for IPv6 firewall filters with match conditions on address prefixes longer than 64 bits, in some corner cases, the filter may not be correctly evaluated and packet loss may occur. [PR879829](#)
- After committing some configuration changes (e.g. deactivate an interface), while the Packet Forwarding Engine daemon (pfed) tries to get statistics of some nodes, it may encounter a NULL node, causing pfed to crash and dump core. [PR897857](#)

### ***General Routing***

- MPC might crash with core-dumps due to watchdog timer expired. [PR593444](#)
- The knob route-memory-enhanced(hierarchy: set chassis) is hidden in platforms M320 and MX series. There is no functionality break but this knob should not be hidden. [PR690100](#)
- For an IPv4 pool, only the all-0 host and the all-1 host addresses are precluded from allocation, both for gateway-assigned and external address assignment. [PR729144](#)
- The next-hop-group knob is not supported under routing-instance hierarchy, but this knob is present under this hierarchy. This PR is opened to remove next-hop-group knob from routing-instance hierarchy. [PR731264](#)
- Reconfiguring a deleted interface with BFD sessions can take up to 20 minutes for the BFD sessions to initialize. [PR786907](#)
- When gr- interface is disabled, the DECAP-NH also needs to be deleted / set to discard. [PR791277](#)
- The ingress family feature unicast Reverse Path Forwarding (uRPF) check execution order was invalidated when Filter Based Forwarding (FBF) was enabled on MX Series

routers with MPCs or MICs. This solution repositions uRPF just prior to Filter Based Forwarding (FBF), so that both actions are compatible and applicable. This applies to both IPv4 and IPv6. [PR805599](#)

- When the 10x10GE PIC (PD-5-10XGE-SFPP) is configured to run in linerate-mode under [set chassis fpc fpc-number pic pic-number] hierarchy, and an input-scheduler-map with Class of Service (CoS) queues including any of queue 4 to queue 7 is applied to an interface on the 10x10GE PIC, the ingress queues may not map correctly to the internal hardware ingress queues. As a result, packet drops may be seen in a higher priority queue than that which is expected. [PR818605](#)
- In subscriber management environment, with dynamic-profile configured for subscribers, with high churn rate of subscribers, memory leak is observed in authd process. This was observed from a login/logout or flapping of 1000 subscribers every three minutes. [PR835204](#)
- Restarting the FPC can terminate the DFWD process and create a core file. This will require a restart of the OpenFlow daemon for the OpenFlow functionality to work properly again. [PR842923](#)
- Aggregate bundle interface with IPv6 Interface stuck in tentative state. Trigger was deactivation/activation of ae-interface. [PR844177](#)
- When the router runs at full scale for a very long period of time, during which it experiences network failures, all SDB logical unit numbers appear to be used up. The lack of unit numbers causes login failures for subsequent additional subscribers. [PR855181](#)
- When an MPC or DPC fails in a specific manner, while failing it continues to send traffic into the switching fabric for a time. With SCBE the fabric ASICs report errors such as these with large counts: chassisd[82936]: %DAEMON-3: New CRC errors found on xfchip 0 plane 0 subport 16 xfport 4 new\_count 17651 aggr\_count 17651 chassisd[82936]: %DAEMON-3: New CRC errors found on xfchip 0 plane 0 subport 17 xfport 4 new\_count 17249 aggr\_count 17249 chassisd[82936]: %DAEMON-3: New CRC errors found on xfchip 0 plane 0 subport 18 xfport 4 new\_count 65535 aggr\_count 65535 With SCB the fabric ASICs report errors such as these (no CRC count is given in this case): chassisd[1486]: CHASSISD\_FASIC\_HSL\_LINK\_ERROR: Fchip (CB 0, ID 0): link 60 failed because of crc errors chassisd[1486]: CHASSISD\_FASIC\_HSL\_LINK\_ERROR: Fchip (CB 0, ID 0): link 61 failed because of crc errors chassisd[1486]: CHASSISD\_FASIC\_HSL\_LINK\_ERROR: Fchip (CB 0, ID 0): link 62 failed because of crc errors chassisd[1486]: CHASSISD\_FASIC\_HSL\_LINK\_ERROR: Fchip (CB 0, ID 0): link 63 failed because of crc errors This can cause DPC(s) to stall and not send traffic into the switching fabric to other DPCs or MPCs. Messages such as these may be reported by the affected DPC(s) : [Err] ICHIP(1)\_REG\_ERR:packet checksum error in output fab\_stream 4 pfe\_id 64 [Err] ICHIP(1)\_REG\_ERR:packet checksum error in output fab\_stream 6 pfe\_id 64 [Err] ICHIP(1)\_REG\_ERR:packet checksum error in output fab\_stream 8 pfe\_id 64 As well as the above messages, messages such as the following may be reported from the affected DPC(s), (the "phy\_q" number must be in the range 256 to 351 inclusive and the "stream" number must be in the range 64 to 111 inclusive). fpc0 imq\_q\_waiting\_queue\_empty: phy\_q 258, telapsed:200ms, mu0:4248, mu1:4248, mu2:4248, loop:40000 fpc0 imq\_q\_disable\_queue:failed, phy\_q:258 fpc0 ICHIP(0) imq\_stream\_disable\_stream() failed to disable physical queue 258 in stream 65 fpc1

imq\_q\_waiting\_queue\_empty: phy\_q 264, telapsed:200ms, mu0:33386, mu1:33386, mu2:33386, loop:43987 fpc1 imq\_q\_disable\_queue:failed, phy\_q:264 fpc1 ICHIP(2) imq\_stream\_disable\_stream() failed to disable physical queue 264 in stream 68 This failure on the affected DPCs persists, and will likely affect all traffic destined to the fabric from affected DPCs. The only temporary resolution is to restart the affected DPCs, which will resume fabric traffic from the affected DPCs. [PR856560](#)

- When the fxp0 interface on a k2re is administratively disabled, the local end shows the link as down while the far end device displays the status as up. [PR862952](#)
- If a router receives the BGP keepalive at time t, the next keepalive is expected at time t+30 secs (+/- 20% jitter). However, right around the time when the next keepalive is expected to be received, the BGP keepalive packet is dropped due to some network issue (e.g. uplink towards peer flaps). During this scenario, retransmission of BGP keepalive message on BGP peer would take long time and the BGP session will be terminated due to hold timer expiry. [PR865880](#)
- In subscriber management environment, with scaling subscribers login (110K DHCP and 20K PPPoE), after restarting one of the line cards which has subscribers, autoconf process might crash and generate core due to memory corruption or memory double free. [PR870661](#)
- The lldpd process might crash if there are multiple unknown type, length, and value (TLV) elements included in received LLDP PDUs. [PR882778](#)
- Rpd might crash when deactivate rib-groups (inet and inet6) under protocols IS-IS, also these rib-groups applied under interface-routes. The core files could be seen by executing CLI command "show system core-dumps". [PR885679](#)
- Observed a traffic-drd daemon might hang once after logging into service PIC and restarting the net-monitor daemon. [PR889982](#)
- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX Series with Dense Port Concentrator (DPC) or Flexible PIC Concentrator (FPC), the feature of Circuit Cross Connect (CCC) and Virtual Private LAN Service (VPLS) are configured, and a ingress filter attached to routing table. When the link on which CCC is configured flapping, DPC or FPC memory leak occurs. [PR908928](#)
- After FPC/MPC is reset or while PPPoA customer login, in rare case, the ppp daemon (jpppd) might get an incorrect value from device control daemon (dcd) which might cause all the new Link Control Protocol (LCP) messages to be ignored and results in static PPPoA sessions can not come up. This problem is seen on MX Series platform products so far, but the problem is mostly common and if other products are using the same version of Junos OS software it might apply to them. [PR912496](#)
- When the NSR switchover happened immediately after a lot of vrf routing-instances were being deleted, garbage lsi interfaces will remain in kernel, while they are removed from RPD. Those garbage interfaces will result KRT queue stack issue upon later lsi re-configuration. [PR912861](#)
- The problem is related to composite next-hops (cnh). Some of the next-hops of that type might end up being not allocated while processing export policy configured under routing-options forwarding-table hierarchy. This issue is particularly visible in a scenarios where cnhs form so called unilist (which usually is a result of ECMP scenario). [PR919489](#)

- An FPC crash can be triggered by an SBE event after accessing a protected memory region, as indicated in the following log: "System Exception: Illegal data access to protected memory!" The DDR memory monitors SBEs and reports the errors as they are encountered. After the syslog indicates a corrupted address, the scrubbing logic tries to scrub that location by reading and flushing out the 32-byte cache line containing that location in an attempt to update that memory location with correct data. If that memory location is read-only, it causes an illegal access to protected memory exception, as reported, and resets the FPC. The above-mentioned scrubbing logic is not needed because even if SBE is detected, the data is already corrected by the DDR and the CPU has a good copy of the data to continue its execution path. PR/919681 can be triggered on both PTX and T4000 platforms and can be seen in Junos OS releases 12.1 and 12.3. Fix is available in 12.3R5, 12.3R3-S6, 13.3R1, 13.2R2. Crash signature in the FPC shell shows the following: SNGFPC4(router-re0 vty)# sh nvram System NVRAM : 32751 available bytes, 2477 used, 30274 free Contents: [LOG] Set the IP IRI for table #1 to 0x80000014 [LOG] IPV4 Init: Set the IP IRI to 0x80000014 [LOG] GN2405: JSPEC V 1.0 Module Init. <..> Reset reason (0x84): Software initiated reset, LEVEL2 WATCHDOG [Sep 6 17:16:07.231 LOG: Warning] <164>DDR: detected 3 SDRAM single-bit errors [Sep 6 17:16:07.231 LOG: Warning] <164>DDR: last error at addr 0x108d2378, bad data/mask0x00240401ffffff7/0x0000000000000008 bad ecc/mask=0xbe/0x00 System Exception: Illegal data access to protected memory! <<< Event occurred at: Sep 6 17:16:07.231087. [PR919681](#)
- Leak in /mfs/var/sdb/iflstatsDB.db. [PR924761](#)
- If rpd ACK feature is enabled through command "indirect-next-hop-change-acknowledgements". When a route being added and a quickly route change happens on the same route, high routing protocol process (rpd) CPU utilization might be seen and stays high (above 90%) until rpd is restarted. [PR925813](#)
- When P2MP LSP is protected by link protection, it could have active and multiple standby next-hops. If one of the next-hops, regardless of whether it is an active or standby one, is removed due to FPC power-off or failure, multicast diagnostics daemon (mcdiagd) falls into infinite loop while collecting next-hop information. [PR931380](#)
- Added AI-Scripts workaround for Junos OS bug sw-ui-misc/920478 (FIPS crash). [PR932644](#)
- In Junos OS versions later than 11.2 where logical interface localization is enabled, Routing Engine mastership switchover could lead to logical interface indexes inconsistency in Ichip FPCs when graceful Routing Engine switchover (GRES) is configured. This inconsistency could gradually lead to logical interface index overlaps and traffic blackholing. [PR940122](#)
- When nonstop active routing (NSR) is configured and the memory utilization of rpd process on the backup Routing Engine is high (1.4G or above), the rpd crash on backup Routing Engine may bounce the BGP sessions on the master Routing Engine. [PR942981](#)
- Egress multicast statistics displays incorrectly after flapping of ae member links on M320 or T Series FPC (M320 non-E3 FPC and T Series non-ES FPC). [PR946760](#)
- When a router is booted with AE having per-unit-scheduler configuration and hosted on an EQ DPC, AE as well as its children get default traffic control profile on its control



logical interfaces. However, if a non-AE GE interface is created on the DPC with per-unit-scheduler configuration, it will get default scheduler map on its control logical interfaces. [PR946927](#)

- When configuring "no-readvertise" flag to existing static route, then this static route will not be exported to other VPN routing and forwarding (VRF) tables from onwards which is expected. However, for the static route has already exported to other VRF tables before "no-readvertise" configuration, no deletion event occurs. Also, the "rt-export" bit still set for the static route which is exported to other routing tables after "no-readvertise" configuration. [PR950994](#)
- Under particular scenarios, commit action might lead the Context-Identifier to be ignored when OSPF protocol refreshes its database. Then the PE router will stop advertising this Context-Identifier. [PR954033](#)
- "show interfaces et-x/y/z extensive" will display MRU now. MRU can be configured at "set interfaces et-x/y/z gigether-options mru". If MRU is not configured, then it is defaulted to MTU + 8. MRU displayed from the CLI does not include the CRC. [PR958162](#)
- For MXVC platform, the pfe reconnect timer extends from the default 15 seconds to 60 seconds temporarily. This will be reversed once Packet Forwarding Engine connection issues are resolved. [PR963576](#)
- On T1600/T640/T320 with FPC installed or M320 with FPC type 1/2/3(non E3 FPC). In processing for fpc-resync and fab-liveness packets (a kind of periodic probe packet) if error occurs while sending packet, the buffer of packet does not be freed. This causes packets buffers to leak and eventually the packet heap runs out of memory. [PR973892](#)
- PPP over ATM transit traffic was not being fragmented correctly by ATM MIC. The changes allow the fragmentation of the transit traffic to work properly. [PR976508](#)
- A raw IP packet with invalid Memory Buffer (mbuf) length may trigger a kernel crash. The invalid mbuf length might be set by other daemons incorrectly. [PR1006320](#)

### **High Availability (HA) and Resiliency**

- On TX Matrix routers with four LCCs and IQ2 PICs, in-service software upgrade (ISSU) from 12.3R1.7 to a newer release results in traffic loss and a FRU upgrade error. [PR768502](#)
- On TX or TXP Line Card Chassis (LCC) with graceful Routing Engine switchover (GRES) enabled, if a mastership switch is being requested on a LCC who's backup Routing Engine (RE)'s em0 interface is physically failed (due to hardware failure or driver stops working), this will cause all FPCs on the LCC to disconnect from the old master Routing Engine, but can not reconnect to the new master one either. [PR799628](#)
- RPD on the backup Routing Engine may crash when it receives a malformed message from the master. This can occur at high scale with non-stop routing enabled when a large flood of updates are being sent to the backup. There is no workaround to avoid the problem, but it is rare and backup rpd will restart and the system will recover without intervention. [PR830057](#)
- Backup Routing Engine sends Arp 128.0.0.6 to Packet Forwarding Engine, then they are counted as "unknown" on show pfe statistics traffic. [PR830661](#)

- PR 855661 will affect IQ2 PICs during unified ISSU on TX platform. When upgrading to 12.3R2 from releases prior to 12.3R2 through unified ISSU, IQ2 PICs will report error. This error is due to IQ2 PICs not being able to download image during unified ISSU. [PR855661](#)
- During a router hardware upgrade procedure, in dual Routing Engines system, the newly installed Routing Engine may overwrite the other Routing Engine configuration with the factory default configuration. As a result, both Routing Engines may bootup in "Amnesiac" mode. This situation can occur under following conditions: - RE0 has default factory configuration and, - RE1 has "commit synchronize" enabled - Both RE0 and RE1 boot-up simultaneously, or - RE0 is UP and running and RE1 is restarted. [PR909692](#)
- If NSR Routing Engine switchover his done right after committing the configuration change which deletes routing-instance(s), some of those instances will not be deleted from forwarding table. [PR914878](#)
- "When LACP is configured in 'periodic fast' mode, the traffic loss of more than 30 seconds will be seen during unified ISSU. The workaround is to change LACP to 'periodic slow' mode before unified ISSU." [PR1059250](#)

### ***Infrastructure***

- On all M/MX/T Series devices, when you log in to the device, the login process might crash due to abnormal disconnection behaviors. [PR802169](#)
- Unsolicited Neighbor Advertisement is not sent from backup when vrrp switchover is initiated. [PR824465](#)
- Kernel log messages may be logged roughly about every 10 minutes - "%KERN-6: MTU for xxxx::yy reduced to 1500" when there is a pmtu reduced to a destination. These are information messages and there is no operation impact due to these messages. [PR888842](#)
- On RE-S-1800 family of Routing Engine, after an intensive writing to SSD, the immediate rebooting might cause SSD to corrupt. [PR937774](#)

### ***Interfaces and Chassis***

- On logical tunnel (lt) interfaces, you might not be able to use the 'family vpls' option at the [edit interfaces lt-fpc/pic/port unit logical-unit-number] hierarchy level. [PR44358](#)
- 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. [PR65800](#)
- In a SAToP pseudowire on a 4-port COC3/CSTM1 or 12-port T1/E1/J1 CE PIC, when there is data loss from the pseudowire or due to an alarm condition (LOS/LOF/AIS) at the peer end of the SAToP pseudowire, the local PIC does not transmit AIS. [PR602563](#)
- When you have the below config on a logical interface, unit 2000 { encapsulation vlan-bridge; vlan-tags outer 40 inner-list [ 20 3000 ]; family bridge; } And you execute "show interface intf-name extensive" you will see the below: Under " Flags: SNMP-Traps Redundancy-Device 0x20004000 VLAN-Tag [ 0x8100.40 0x8100.2000 20,3000 ]

", you will see the unit number 2000 between outer and inner tags configured. This is just a display issue and no functionality is affected. [PR723188](#)

- To troubleshoot a particular subscriber, one can use 'monitor traffic interface <ifd> write-file xy.pcap'. Using this command on aggregated or demux interfaces can lead to corrupted ingress packets in the PCAP file. Customer traffic is not affected though. [PR771447](#)
- Collecting subscriber management control traffic via 'monitor traffic interface demux0 write-file xy.pcap', the logical unit number is incorrect when multiple demux IFL's are present. This problem is fixed and the correct interface logical unit number is reported in the juniper header of the captured PCAP file. [PR771453](#)
- In MLPPP scenario, in rare conditions (such as FPC crash), kernel may try to delete a MLPPP bundle with an invalid (although within the max bundle limit) bundle ID. This will casue vmcore and Routing Engine switchover. [PR780784](#)
- On an MX Series system, if a composite next hop entry which is deleted in Master NW-INE exists in Backup NW-INE, the kernel may crash with a vmcore file generated during NW-INE switchover. [PR793098](#)
- Master LED of craft interface keeps Green during Halt the system or Power off. [PR805213](#)
- With LSQ interface, the MLPPP fragments cannot use the egress queue 4 to 7 on the MLPPP member links. There is no workaround. [PR805307](#)
- Junos OS does not generate VRRPv4 mastership change syslogs while it generates VRRPv6 logs. [PR807217](#)
- LFM action profiles may execute multiple times if 1. more than one event tlv is present in received event pdu, or 2. link adjacency loss happens after seeing a threshold crossing event tlv. It may be noted that this redundant execution of action profile is not feature impacting. [PR816153](#)
- With Junos OS 11.4 or higher and Enhanced SCB installed on a mix of MX Series and DPC cards REG\_ERR messages might be reported under certain traffic flow conditions from MX Series to DPC card. On the receiving DPC card fabric cell received out of order will be re-ordered and merged to build the packet. If this out-of-order delivery is too high a reorder event will be triggered and all cells belonging to the packets are dropped. The frequency is low rate. The following syslog entry will be reported Sep 29 20:43:10 node fpc8 ICHIP(3)\_REG\_ERR:first cell drops in ichip fi rord : 4122 Sep 29 20:43:10 node fpc8 ICHIP(3)\_REG\_ERR:Non first cell drops in ichip fi rord: 7910. [PR821742](#)
- After the corresponding FRU reset, the FPC or SIB alarm still did not get cleared, such as dest error , sib link error and sib check. Alarm time Class Description 2012-11-23 18:56:09 EST Minor FPC 0 dest error 2012-11-23 18:56:09 EST Minor FPC 6 SIB Link Error 2012-11-23 18:56:08 EST Minor Check SIB 4. [PR836830](#)
- If the "tunnel-destination" address of a Generic Routing Encapsulation (GRE) interface is placed in one instance and the GRE interface is placed in another routing-instance, the lookup for the GRE tunnel destination is done on inet.0 instead of the appropriate routing instance's inet.0 table. The similar issue could happen on IP-over-IP or Automatic Multicast Tunneling (AMT) tunnels too. [PR851165](#)

- The device configuration daemon (dcd) may crash when a partial demux subinterface configuration is attempted to be committed. There is no impact to traffic forwarding but before the configuration can be committed, it must provide a valid 'underlying-interface' for the demux subinterface. [PR852162](#)
- On M/MX/T Series platform with Services PIC and dual Routing Engines, configure MPLS and set Services PIC in layer-2 mode. Apply Class-of-Service (CoS) configuration to sp-x/y/z control interface. After that perform graceful Routing Engine switchover (GRES), then the Services PIC might restart. [PR859036](#)
- MC-LAG will no longer change just the LACP System Identifiers directly, but will also remove the "Synchronization, Collecting, Distributing" bits from the Actor State bits advertised in the PDU. [PR871933](#)
- Memory leak in PPPoE daemon is seen while running repeated concurrent login /logout of PPPoE subscribers. [PR874006](#)
- On MX Series routers with MPC with 20port GE MIC, interface stores packets when disabled and transmits stored packets after enabled. [PR874027](#)
- "Link down" alarms should never exist on the VC Protocol Backup Routing Engine. They should only be on Protocol Master, if any. The bug is that the "Link down" alarms are not cleared from the Protocol Backup after/during a GRES event. Restarting alarmd removes these alarms from the Protocol Backup. [PR886080](#)
- To configure FEC thresholds via CLI, use string format with mantissa and exponent: Example: set interfaces et-1/0/0 otn-options signal-degrade ber-threshold-signal-degrade 1.23E-4 set interfaces et-1/0/0 otn-options signal-degrade ber-threshold-clear 2.34E-5. [PR886572](#)
- The C-LMI (Consortium LMI) is supported on all i-chip based FPC. Support for the MX-FPC 2 and 3 was missing and now added. [PR895004](#)
- In current MX Series implementation, the physical or logical interfaces (ifd/ifl) might be created and marked UP before a resetting FPCs' fabric planes are brought up and ready to forward traffic, as the result, traffic might be blackholed during the time window. This window of traffic blackhole is particular long if the chassis is heavily populated with linecards, for example, the router has large scale of configuration (routes or subscribers), and coupled with a lot of FPC reset, such as upon a node power up/reset. [PR918324](#)
- Non-Existent leg in AE bundle prevents DHCP subscribers from coming up. [PR918745](#)
- Queue stats counters for AE interface will become invalid after deactivating ifl on the AE interface. [PR926617](#)
- PPPoA session would not come up on removal/addition of cable to the tester port. [PR939404](#)
- Strange FRU Insertion trap [RE PCMCIA card 0] is generated when Routing Engine master-switching is done on box with RE-1800. [PR943767](#)
- When transit traffic of Ethernet frames of size less than 64 bytes are received by 1x 10GE(LAN/WAN) IQ2E PIC, the router forwards the frames instead of dropping them. [PR954996](#)

- When a logical interface containing some vrrp group configuration is deleted, snmp walk on vrrp MIB may loop continuously. [PR957975](#)
- In very uncommon situation, we will see LCCs chassisd state is inconsistent with SFC chassisd state, this is very misleading in troubleshooting stage. This PR fixed this issue. [PR963342](#)
- In the multilink frame relay (mlfr) environment with "disable-tx" configuration. When the differential delay exceeds the red limit, the transmission is disabled on the bundle link. When it is restored, the link should be added back. But in this case, the link stays disable state and it is not rejoined to the bundle. [PR978855](#)
- In Ethernet OAM connectivity-fault-management, Junos OS default encodes MAID(MD name and MA name) in character format. Currently only 43 octets are supported in Junos OS for the MD + MA name. Junos OS needs to support maximum length of 44 octets for MAID per the standards. [PR997834](#)
- On MX Series router with MX Series linecard or T4000 router with type5 FPC, when the "Hardware-assisted-timestamping" is enabled, the MPC modules might crash with a core file generated., The core files could be seen by executing CLI command "show system core-dumps". [PR999392](#)

### **Layer 2 Features**

- When directly apply sampling on VPLS interface (i.e interface ge-4/0/1 unit 0 family vpls sampling input), if customer configures logical interface and sampling input/output together first time, then deactivating sampling input/output through CLI, kernel will then not disable the sampling. Also note that, the action of sampling is a hidden command for VPLS interfaces and would not be listed in "possible completion" list when combined with "?". [PR772270](#)
- On MX Series systems after the changes performed within PR/686399 Junos 10.4R9 or higher, traffic destined towards mac addresses learned from the core interfaces are aged out every aging interval and added again. During this very short event, VPLS traffic will get flooded. [PR820726](#)

### **Layer 2 Ethernet Services**

- Traffic loss after performing graceful Routing Engine switchover (GRES). Two similar problems are fixed here: 1. In the rare case, after the first GRES, some IPv6 routes fail to be added because the buffer to routing protocol daemon (rpd) is full and thus the response to the add request fails. 2. Somewhere, when one GRES is performed after another GRES, some IPv4 routes fail to be added because the logical interface is not up yet and the Next-hop address isn't populated in a timely manner. [PR808932](#)
- To free the socket used by jdncpd for bootp helper, deactivate dhcp-service traceoptions. [PR817515](#)
- "show bridge mac-table interface X vlan-id Y" is empty on trunk port. This is just a display issue. This MAC is present on the forwarding table that can be confirmed using command "show route forwarding-table family bridge". [PR873053](#)
- If STP is configured on AE interface, the l2cpd might be under high utilization and VRRP repeatedly flaps after the VRRP active router reboots. The root cause here is when

STP is configured on AE interface, the corresponding Bridge Protocol Data Unit (BPDU) messages will go to Routing Engine instead of processed in Packet Forwarding Engine. [PR882281](#)

- When toggling VLAN tagging type from "flexible-vlan-tagging" to "vlan-tagging" or vice versa, the integrated bridging and routing (IRB) MTU should be changed accordingly. However the IRB MTU is not re-computed in this case, which might lead to connectivity outage. [PR928746](#)
- In DHCPv6 subscriber environment, changing the c-tags (inner vlan) without clear the DHCPv6 clients first is not recommended, it might cause the subscriber to use the old inner vlan even after DHCPv6 RENEW process. [PR970451](#)

### **MPLS**

- For point-to-multipoint LSPs configured for VPLS, the "ping mpls" command reports 100 percent packet loss even though the VPLS connection is active. [PR287990](#)
- Statistics for a p2mp lsp used for CCC connection will not be displayed on a Ingress PE. [PR444336](#)
- MPLS forwarding broken when labeled BGP routes are distributed to LDP. [PR724658](#)
- A-----B-----C X(primary) X(protector) Y(protector) Y(primary) Given a topology with primary and protector context IDs like above, B seems to have a problem installing LDP label routes in mpls.0 and the forwarding table. [PR782499](#)
- In an RSVP environment with AutoBw, the Bandwidth Adjustment timer for new LSPs added simultaneously is not smeared along with the rest of the existent LSPs when the smearing algorithm is triggered. [PR874272](#)
- In a scenario with scaled MPLS tag labels exist, while MPLS flapping (which could be triggered by routing protocol flapping), routing protocol daemon (rpd) might crash and dump core due to system tries to delete an already freed MPLS tag label Element. [PR878443](#)
- The rpd might crash in specific RSVP-Signaled Point-to-Multipoint (P2MP) LSP re-merge scenario. [PR890787](#)
- In current Junos OS, lsping/lsptrace utils have compatibility issue with other vendor routers. millisecond field might show huge value which result incorrect RTD calculated. Juniper-MX960> ping mpls ldp 192.168.228.7/32 source 192.168.199.193/32 exp 5 count 5 size 100 detail Request for seq 1, to interface 510, label 1102, packet size 100 Reply for seq 1, return code: Egress-ok, time: 3993729.963 ms <--- Local transmit time: 2013-04-29 12:05:06 IST 873.491 ms Remote receive time: 2013-04-29 12:05:06 IST 3994603.454 <---- This is cosmetic issue and current software limitation. [PR891734](#)
- LSP metric will be not correctly changed as the new configured one after committed when cspf finds an Explicit Route Object (ERO) different from the current ERO and the Path State Block (PSB) re-signaling fails. This is because a change in metric is a local PSB change, but after a configuration change (for example, the bandwidth requirement was changed), PSB and associated routes used to get this change only after a cspf computation followed by a session refresh or re-signaling. If the re-signaling

fails, the configured metric value is not updated in the existing PSB and the route metric. [PR894035](#)

- RPD might crash after executing "ping mpls l2vpn interface <interface>" command under specific time window. [PR899949](#)
- The output of "show ldp overview" command regarding graceful restart is based on per protocol LDP graceful restart settings. Where graceful restart is enabled by default. So when graceful restart is disabled this command shows it's enabled for LDP. However graceful restart should be enabled globally for LDP graceful restart to operate. [PR933171](#)
- The RSVP bandwidth of the aggregated ethernet (AE) bundle does not adjust properly when a member link is added to AE interface, and at the same time an IP address is removed from this AE bundle. [PR948690](#)
- MPLS traceroute causes "rttable-mismatch" syslog messages. [PR960493](#)
- During SNMP walk on table MPLS cross-connect table (mplsXCTable) in case of flood nexthop, the rpd might crash. [PR964600](#)
- In the MPLS environment, when execute the command "show snmp mib walk mplsXCTable" to walk the MPLS cross connect table, the routing protocol daemon (rpd) CPU utilization might reach over 90%, and the rpd process doesn't respond to any CLI show commands. [PR978381](#)
- LSP metric modification leads to Constrained Shortest Path First (CSPF) computation and resignaling. It should update RSVP routes directly. [PR985099](#)
- In the MPLS environment with "egress-protection" configuration, there is a direct LDP session between primary PE and protector. One context-id is configured as primary PE's loopback address or any LDP enabled interface address. When delete the whole apply-group or delete the ldp policy from apply-group, the routing protocol daemon (rpd) might crash. [PR988775](#)
- In l2circuit scenario with LDP session established between Juniper Networks PE and Cisco PE, if Cisco PE is not sending a label withdraw for the l2circuit Forwarding Equivalence Class (FEC) before advertising a new label for it, and later, when Cisco PE tries to change the l2circuit parameters, the rpd process might crash on Juniper PE. This issue does not occur in Junos OS only environment as it always sends a label withdraw before advertisement of new label. [PR1016270](#)

### **Network Management and Monitoring**

- Any 'show snmp mib walk' cli command, in case of getting in a loop, will introduce SNMPD memory leak, then eventually crash. Note: The same snmp walk from remote NMS won't trigger this issue. [PR732852](#)
- Source/destination values in SNMP Get-Response messages are not correct in the SNMP log file. These values should be exchanged. But actually SNMP Get-Response messages are sent properly with correct source and destination addresses. [PR784780](#)
- Junos OS version later than 10.0 reserves separate index pool for private and public interfaces. After an upgrade from version prior to 10.0 to version later than 10.0, some of the public interfaces may be included within the private index pool. There is no operational risk caused by this issue. [PR815028](#)



- When OID dot3adAggPortTable is polled on the router, SNMPD tries to fetch interface/interface-unit data from kernel in ASYNC mode. It is possible that by the time kernel replies with stats, the interface/interface-unit state is already changed or deleted. This problem is most commonly seen when kernel reply for interface/interface-unit are late, more than expected window. Accessing interface/interface-unit stats for which the state has already been changed can cause the MIB2D core-file. At such times MIB2D must first check if the interface/interface-unit entries are still valid (i.e. not changed or deleted) before accessing the associated data-structure. A check has been added via this PR which marks the interface/interface-unit as stale if state has been changed, this prevents the code from accessing any associated data-structure if the entry is marked as stale. [PR852282](#)
- Mib2d may get ATM VPI updates before the ATM IFDs are learnt. In such cases instead of discarding the updates, mib2d has started caching them until the IFD is learnt. [PR857363](#)
- When we do SNMP polling via CLI on a big MIB node which has lots of OIDs and huge data, like "show snmp mib walk 1.3.6.1.4.1". CLI might not be able to consume data at the rate it was being generated by snmpd, so the snmpd buffer is occupied more and more, eventually this would cause snmpd to reach its limit then crash. [PR864704](#)
- SNMP query from valid client on routing-instance-1 with community string that belongs to routing-instance-2 gets the details of routing-instance-2 instead of blocking such queries based on community. [PR865023](#)
- While some set operation is in progress there is a huge pile up of pending requests in netsnmp\_agent\_queued\_list Queue., which is running into several thousands of requests which is causing the memory consumption to increase in snmpd and running out of 256 MB of rlimit and crashing. [PR920471](#)
- When syslog server is configured using hostname, after Routing Engine switchover router stopped sending the syslogs to external syslog server. Immediately after switchover, DNS was not accessible because it will take some time to learn route to DNS. System stopped retrying DNS resolution and syslogging stopped. System was running GRES (no NSR). [PR947869](#)

### ***Platform and Infrastructure***

- Commit time warning is changed to trace message. [PR480082](#)
- On the process details page (Monitor > System View > Process Details) of the J-Web interface, there are multiple entries listed for a few processes that do not impact any functionality. [PR661704](#)
- High CPU utilization (100%) will be seen on MPC if inline sampling/Jflow is enabled. The high CPU will affect the MPC's normal tasks. [PR671136](#)
- On certain M Series routers (M20, M40, M40e, and M7i/M10i without Enhanced CFEB), the following error message is displayed on the Packet Forwarding Engine console periodically: "pfe\_get\_ifl\_stats failed" This error is seen only if aggregate interfaces (like AE or AS) are configured on the router. There is no functional impact because of this error message. [PR692081](#)



- On MX Series router with MPC, the traffic which needs across GRE tunnel is still forwarded after disabling the GRE interface. [PR707140](#)
- A rare race condition may occur when updating the forwarding table for a VPLS instance that is using no-tunnel-services with an IRB interface and igmp-snooping. The MPC will crash and restart as a result, leaving a core file behind. [PR741999](#)
- In l3vpn setup, customer facing interface fails to forward traffic, if RPF and localization is enabled in sequence. [PR752540](#)
- On the JCS-1200 RE-JCS-1X2400-48G-S Routing Engine configuration of the MAC address on the external interfaces em0 and em1 is not allowed. You cannot configure the MAC address on fxp0 on the other Routing Engines supported on the JCS-1200 as well. Therefore, the Junos CLI to configure the MAC address on em0 and em1 interfaces has been disabled. [PR770899](#)
- The unit for IPv6 router Advertisement "Reachable time" and "Retrans timer" presented on the CLI (monitor traffic interface) is in second which is incorrect. The unit for these timers should be in milliseconds. Tcpdump utility had a bug in decoding icmp6 parameters, such as reachable-time and retransmit-time. It was just a display issue. The unit for "router-lifetime" presented on the CLI is in second which is correct. [PR796672](#)
- NPC core observed @nh\_dfw\_get\_correct\_next\_intf\_prefix. [PR801607](#)
- When changing configuration repeatedly, in rare conditions, some internal errors may cause CLI process hogs memory and the utilization keeps on increasing due to memory leak. When the memory usage of CLI process increases to around 85% of system limit, the following logs could be seen: /kernel: Process (1383,cli) has exceeded 85% of RLIMIT\_DATA: used 62048 KB Max 65536 KB The memory will be released once user logout from the router. [PR813673](#)
- Commit may fail, when a config object is deleted and re-added as transient change from a commit script. [PR814796](#)
- CLI command 'show route forwarding-table' would only display <= 16 ecmp paths when CBF is used. [PR832999](#)
- An FPC may reboot when a live-core is requested and the /var partition does not have sufficient space to store the live-core. [PR835047](#)
- IPv6 traceroute is not setting traffic class with TOS command. Traceroute packets may not get to the correct queue and the COS bits may not be reflected properly in the traceroute outputs. [PR835359](#)
- Added support for "raise-rdi-on-rei" knob on FPCs on MX Series and T Series routers. [PR844097](#)
- The audit daemon (auditd) is the daemon which handles system accounting events and tries to send them out to configured RADIUS servers. If there is any problem in sending these accounting records to RADIUS (In this case RADIUS servers are unreachable or disconnecting frequently), auditd will spend more time on each accounting record because of the retries, and during this time if there are many accounting events, all those records will be in queue. And at one point of time, queue exceeds its limit and hence auditd crashes. [PR863697](#)

- In the Network Time Protocol (NTP) configuration, if the specified source ip address is not in current routing-instance, the router will use primary address of interface (which will be used to send packet) as source address. Client routers will treat the NTP packets as incorrect packets, and then NTP synchronization failed. [PR872609](#)
- After multiple iterations of active FPCs restart and GRES, E2-FPC crash because of hogging CPU. [PR873718](#)
- When we are deleting a configuration hierarchy which has no groups applied, the corresponding group object hierarchy is also marked as changed in commit script view. [PR878940](#)
- When the instance has vlan-id all and adding interface unit with "vlan-tags outer X inner Y" to this instance, traffic from ALL instance VLANs is leaking over that unit tagged with outer tag X and each VLANs own inner tag A,B,C,..... When the instance has vlan-id all, for dual tagged ifl, the inner vlan check will be done. [PR883760](#)
- In DHCP relay agent scenario, DHCP offers message with option82 (relay-agent-option) is discarded by UDP Forwarding process (fud) after receiving the reply back from DHCP server. This issue happens when the length of interface name (including underlying and parent interface) greater than 23. For example: irb.1011/0/0.1011 - 22 characters works irb.1011/0/0.10011 - 23 characters fails. [PR886463](#)
- On all M/MX Series devices, when a router is acting as an NTP broadcast server, broadcast addresses must be in the default routing instance. NTP messages are not broadcasted when the address is configured in a VPN virtual routing and forwarding (VRF) instance. [PR887646](#)
- On MX Series routers with MPC, firewall filter counter doesn't count packets when firewall is configured on discard interface. [PR900203](#)
- When the ATM interface is configured with hierarchical scheduling, a traffic-control-profile attaches at ifd level and another output traffic-control-profile attaches at ifl level. After the interface restart, the CoS information might not be re-applied correctly, and the packet might be lost. [PR908807](#)
- Changing the domain-name doesn't reflect in DNS query unless a commit full is done. This bug in management daemon (mgd) has been resolved by ensuring mgd propagates the new domain-name to file /var/etc/resolv.conf, so that this can be used for future DNS queries. [PR918552](#)
- With xml:warning and xml:error enabled inside commit scripts, when there is an XML tag mismatch detected in any of the commit scripts, the following errors are seen: error: [filename: xnm:rpc results] [line: 771] [column: 7] [input: routing-engine] Opening and ending tag mismatch: routing-engine line 7 and rpc-reply error: [filename: xnm:rpc results] [line: 773] [column: 6] [input: rpc-reply] Opening and ending tag mismatch: rpc-reply line 6 and junoscript error: [filename: xnm:rpc results] [line: 774] [column: 2] [input: junoscript] Premature end of data in tag junoscript line 2. [PR922915](#)
- DDOS\_PROTOCOL\_VIOLATION alarm shows incorrect timestamps <time-first-detected> and <time-last-detected> on messages. Both fields indicate the same timestamps. Timestamps <time-first-detected> and <time-last-detected> are overwritten. [PR927330](#)

- The `jcs:dampen()` function will not perform correctly if the system clock is moved to an earlier time. [PR930482](#)
- When replacing ichip FPC with MX Series FPC, "traceroute" packets going through an MX Series FPC may experience higher drop probability than when using an Ichip FPC. [PR935682](#)
- The Routing Engine and FPCs are connected with an internal Ethernet switch. In some rare case, the FPCs might receive a malformed packet from the Routing Engine (e.g. packet gets corrupted somewhere on it's way from Routing Engine to FPC), then the toxic traffic might crash the FPC. [PR938578](#)
- FPC might crash with CPU hog due to excessive link flaps causing the interrupts to go high. [PR938956](#)
- On a router which does a MPLS label POP operation (penultimate hop router for example) if the resulting packet (IPv4 or IPv6) is corrupted, then it will be dropped. [PR943382](#)
- On MX Series based line cards, certain ATM packets carried through MPLS backbone might be misidentified as Ethernet in MPLS payload. This misidentification would cause MX Series to believe these packets should be load balanced when they shouldn't. This in turn causes the packets to show up out of order on their ATM network. The out of order packets are causing CRC Input errors, and length errors and the packets to be dropped. [PR946694](#)
- On MX Series based line cards, there is an extremely tiny window that when one interface is created short after another interface, which has packets getting exception handled, get deleted, the created interface might reuse the same interface index as the deleted interface. In this case, the NPC might crash when the exception packets come in before the interface is fully created. [PR960029](#)
- In multi-chassis platform, one of LCC's mastership change causes other LCC's SPARE-SIB's Active-LED to be set abnormally instead of "actual active plane's LED". There is no impact on operation, it is a cosmetic issue. \* only if spare-SIB is SIB#0. For example, - SCC-RE0(M),RE1(B) | LCC0-RE0(M),RE1(B) | LCC1-RE0(M),RE1(B) - all-chassis SIB0 is spare status. - LCC0's mastership change makes the issue on LCC1. - LCC1's spare-SIB0's active LED to be set abnormally. [PR972457](#)
- Under some situations, the MAC entry on a Packet Forwarding Engine is not up-to-date and the frames targeted to a known MAC address will be flooded across the bridge domain. [PR1003525](#)

#### ***Routing Policy and Firewall Filters***

- The auto-complete feature is not working for the "show policy" command. [PR471332](#)
- If RPF and/or SCU is enabled, then any change to an ingress firewall table filter will trigger RPF/SCU reconfiguration for every prefix in the routing table. This may cause transient high CPU utilization on the fpc which may result in SNMP stats request being timed out. [PR777082](#)
- Configuration of an extended community such as: `rt-import:*:* src-as:*:*` fails because the wildcard is not allowed during the configuration validation process. [PR944400](#)

### ***Routing Protocols***

- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has 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. [PR51975](#)
- The problem is encountered when "show route summary table" command is executed with either no table name, or a table name that does not actually exist. When the keyword "summary" was encountered, a background task was started that is used to ultimately display the gathered information. When no table was found to get information from, due to no name given, or incorrect name given, the memory assigned to control the background task was mistakenly freed. When the background task came active, it was accessing memory that had been freed, and ultimately caused the failure. To fix the problem, the code now checks to see if a background task has been started for the command. If there is no background task, the memory is freed as before. If there is a background task, it leaves the memory intact for the background task to use and take care of. The issue was introduced in 12.1 Junos after adding the new feature in the release. Old release should not be affected by the issue. [PR734560](#)
- After upgrade to 10.4R9, the following message are seen "Cancelling deferral pp0 index 131" These messages are not indicative of any problem and only cosmetic. [PR742534](#)
- When the upstream interface for (\*,g) is the same as (s,g)'s, and when it is down or flapping, the SPT-bit on the (s,g) entry is incorrectly cleared. And then multicast forwarding table (mroute) is not updated even after the router receives (\*,g) or (s,g) join from downstream device. [PR753526](#)
- On single Packet Forwarding Engine routers (MX 80 and ACX), PPMD (Periodic Packet Management Daemon) can distribute BFD over AE without installing rules. [PR773101](#)
- With this fix, "jnxBgpM2PrefixesInPrefixesRejected" counter will return the number of prefixes from a BGP peer, that are not eligible to become active. This change makes the variable conform to definition in the specification <http://tools.ietf.org/html/draft-ietf-idr-bgp4-mibv2-03>. There is a new variable "jnxBgpM2PrefixInPrefixesActive" introduced, to return the number of active prefixes from a BGP peer. So the new sequence of variables for the table is as follows:  
root@root> show snmp mib walk jnxBgpM2PrefixCountersTable  
jnxBgpM2PrefixCountersAfi.0.1.1 = 1 jnxBgpM2PrefixCountersSafi.0.1.1 = 1  
jnxBgpM2PrefixInPrefixes.0.1.1 = 0 jnxBgpM2PrefixInPrefixesAccepted.0.1.1 = 0  
jnxBgpM2PrefixInPrefixesRejected.0.1.1 = 0 jnxBgpM2PrefixOutPrefixes.0.1.1 = 3  
jnxBgpM2PrefixInPrefixesActive.0.1.1 = 0. [PR778189](#)
- This issue reported captures a change in behavior observed from previous releases. The adjacency hold down is taking longer than expected on passive interfaces and subsequently the issue disappears. This will not cause any functionality break since the functionality is restored eventually and seen only on passive interfaces immediately after unified ISSU. [PR780684](#)

- In L3VPN scenario, if PE-CE's link is multi-access LAN, the direct subnet route on LAN PE-CE interface will be advertised with a matching nexthop label. In case there are multiple matching nexthops, one of the nexthop labels is selected randomly for the direct subnet route. If the chosen nexthop is unreachable, L3VPN customer's traffic destined to the direct subnet will be dropped. [PR781685](#)
- Routes not deleted from routing table when interface is deleted. Analysis: SPF calculation was not triggered in one particular code flow after the LSAs are deleted from database. Solution: SPF calculation is triggered when the LSA is being deleted due to zero links. [PR782029](#)
- With OSPFv3, PIMv6 or LDP configured, the periodic packet management daemon (ppmd) takes responsibility for these protocols' adjacencies. In a rare condition, kernel might send an invalid packet with a null destination in the message header to ppm process, causing ppm process to crash and dump core. [PR802231](#)
- Continuous soft core-dump may be observed due to bgp-path-selection code. RPD forks a child and the child asserts to produce a core-dump. The problem is with route-ordering. And it is auto-corrected after collecting this soft-assert-coredump, without any impact to traffic/service. [PR815146](#)
- Autoexported secondary BGP routes that are advertised using "advertise-inactive" can miss the flash event; if so, this leaves the deleted secondary route in the stuck state. [PR818552](#)
- OSPF route will not be deleted from routing/forwarding table if configuration satisfies below simultaneously. 1. Router ID is not specified and it can be changed due to interface down. 2. There is an interface where OSPF is not running. Suppose OSPF is running on interface A and it is not running on interface B. IP address of interface A is selected as router ID. When interface A goes down and router ID is changed to the IP address of interface B, OSPF on interface A will lose adjacency to the remote OSPF router but router will keep routes learned via OSPF. [PR820909](#)
- In a rare condition, the periodic packet management process (ppmd) might crash during freeing connections. This problem might be seen if following conditions are met:  
\* BFD is configured \* PPMD connection flag assignment stuck in race condition  
[PR825522](#)
- The case occurs in the BGP multipath scenario. In this case, three EBGP CE's inject same route to VRF (the origin attribute of these routes are "Incomplete"). One of the routes is selected, based on the order the orders come in, and an older one is selected (not the one with lowest router id). After that, two remote PE's also inject same route which is imported as secondary (the origin attribute of these routes are "IGP"), then the remote PE (secondary) routes become active due to origin attribute ("IGP" is preferred over "Incomplete"). When the remote PE (secondary) routes become inactive, a different EBGP CE path becomes active (lowest router-id). The original EBGP CE path that has the multipath nexthop is never turned down and it still retains the multipath nexthop, whereas the new EBGP CE path does not have any multipath nexthop, that leads to no multipath in dataplane. [PR835436](#)
- When inter operate with Cisco router, OSPF adjacency might be brought down by Cisco end, if Junos CPU is high and LSA ACK is delayed for over two minutes. [PR846182](#)

- Junos OS label block allocation can only return block size as power of 2 (e.g. 2, 4, 8, 16,...). In inter-as option-b L2VPN scenario, routing protocol daemon (rpd) core is seen when the ASBR receives a non-power-of-2 label block size from other vendor's device. The root cause here is when rpd requests the non-power-of-2 label block size, an assert occurred. The core files could be seen by executing CLI command "show system core-dumps". [PR848848](#)
- ISIS prefix-export-limit and NSR switchover might push routers into overload when the prefix-export count is clean (=0). [PR853328](#)
- Whenever a configuration change is made and a commit is issued, the Routing Engines CPU utilization could go up due to BGP reprocessing all the routes, because of the commit. This would happen for any commits unrelated to policy, bgp configuration and most common with scaled bgp environment. [PR853670](#)
- When an import-policy change rejects a BGP-route previously contributing to BGP-Multipath formation, the Peer Active-route-counters in "show bgp neighbor" may not get updated correctly. [PR855857](#)
- Multicast packets coming with source address as 0.0.0.0, might cause the rpd to crash. [PR866800](#)
- If the SNMP MIB for BGP is walked, the AFI=1, SAFI=5 entries are missing. If an SNMP "get" is performed, the values can be retrieved. [PR868424](#)
- In VPLS multi-homing environment, with same route-distinguisher configured for the VPLS primary PE and backup PE, routing protocol daemon (rpd) may crash and dump a core file in each of following two scenarios: 1 - On VPLS backup PE, enable "advertise-external" knob, then rpd process crashes and dumps a core file on backup PE. 2 - On VPLS primary PE, enable "advertise-external" knob, after disabling the VPLS interface, rpd process crashes and dumps a core file on primary PE. When issue happens, the following behavior could be observed: user@router> show bgp neighbor error: the routing subsystem is not running user@router> show vpls connections error: the routing subsystem is not running. [PR869013](#)
- In a scenario with graceful restart (GR) enabled for BGP between Cisco platform and Juniper platform, Junos OS is helper (default) and Cisco being restarting router, when Cisco restarts BGP process, Juniper deletes all BGP routes due to doesn't receive End Of RIB (EOR) markers for all configured NLRI's from Cisco. [PR890737](#)
- When the interface goes down, the direct route for that peer address is removed from the routing table before BGP processes interface down event and bring down the session. When BGP calculate multipath routes, since the knob "accept-remote-nexthop knob" is configured, BGP needs to determine whether we can reach the nexthop address (ebgp peer address) directly. BGP did not find direct route for this nexthop address and so asks for route nexthop resolution. In this case, the first BGP path from the peer with up interface has direct router nexthop, the second path is set to have indirect nexthop due to the down interface, BGP passed a wrong mixed multipath nexthop, which caused RPD crash. [PR917428](#)
- If Node-link protection is required in case of multiple ECMP primary paths, Node-link protection command: ("set protocols ospf area <area\_id> interface <interface\_name> node-link-protection") needs to be configured on all the outgoing-interfaces of

PLR(Point of Local Repair)node that fall on the ECMP path to the primary. For eg.in the following diagram: PLR: RTA Destination: RTC Primary paths:

RTA-->lt-1/2/10.102-->RTB-->lt-1/2/10.203-->RTC;

RTA-->lt-1/2/10.122-->RTB-->lt-1/2/10.203-->RTC; Outgoing interfaces on PLR:

lt-1/2/10.102 lt-1/2/10.122 Node-link protection needs to be enabled on both lt-1/2/10.102 and lt-1/2/10.122 if backup route avoiding RTB needs to be computed. (cost 1)

```
|-----|-----lt-1/2/10.102(81.1.2.2 )-----|-----| | (cost 1) | | RTA
|-----lt-1/2/10.122(82.11.22.2)-----| RTB | | | | |lt-1/2/10.203
| 81.3.3.3 | | (cost 1000) |-----| | |lt-1/2/10.103(81.1.3.1) -----| RTC |-----|
|-----| PR924290
```

- In scaled BGP routes environment (global table ~1.5 million routes). First flapping one BGP session (e.g. change the BGP authentication method can get it), after that deleting another BGP session that holds the active routes, this might lead to routing protocol daemon (rpd) scheduler slips. [PR928223](#)
- When nonstop routing (NSR) is configured and path-selection is changed, there might be a non-functional impacting rpd core during the commit process. [PR928753](#)
- Packet Forwarding Engine continuously forwards multicast to Routing Engine as resolve route when multicast is received from non-rpf interfaces. And when the data rate is very high, the Packet Forwarding Engine might be too busy to deal with other good streams. [PR937348](#)
- In Inter-AS FEC-129 VPLS scenario with route-reflector (RR) in each AS, if the RR router is also a VPLS PE, it stops reflect VPLS routes received from MP-EBGP peer to its clients. [PR980834](#)
- In Junos OS, by default the RIP protocol "send" option is set to Multicast RIPv2. When this "send" option is changed from "multicast"(active) to "none"(passive) or vice-versa, rpd core might be seen on the router. [PR986444](#)
- In the P2MP environment with OSPF adjacency are established. One router's time is set to earlier date than another router. OSPF adjacency might not come up when one router goes down and comes up. [PR991540](#)
- When all the following conditions are met, if the knob "path-selection always-compare-med" is configured, the rpd process might crash. - routing-instance (VR, VRF) with no BGP configuration - rib-group in default instance with routing-instance.inet.0 as secondary-rib - rib-group applied to BGP in default instance - BGP routes from master tables (inet.0) leaked to the routing-instance table (routing-instance.inet.0) [PR995586](#)
- When inet.3/inet6.3 is not enabled, BGP group uses inet6.0 table to advertise the routes for both inet6 unicast and inet6 labeled-unicast families. When BGP family is changed, BGP sessions re-establish. When BGP starts to advertise routes to the peer, BGP expects to see route label. However if the old inet6 unicast routes are still present (not completely cleaned), then rpd process crashes. The fix is to separate bgp group for inet6 unicast with inet6 labeled-unicast with same rib. The old peers are cleaned up in the old group and new peers are established in new group. Thus, new peer establishment is not delayed by the cleanup of the old peer. [PR1011034](#)



### ***Services Applications***

- When you specify a standard application at the [edit security idp idp-policy <policy-name> rulebase-ips rule <rule-name> match application] hierarchy level, IDP does not detect the attack on the nonstandard port (for example, junos:ftp on port 85). Whether it is a custom or predefined application, the application name does not matter. IDP simply looks at the protocol and port from the application definition. Only when traffic matches the protocol and port does IDP try to match or detect against the associated attack. [PR477748](#)
- When sending traffic through IPsec tunnels for above 2.5Gbps on an MS-400 PIC, the Service-PIC might bounce due to prolonged flow control. [PR705201](#)
- With EIM and EIF enabled, the internal memory data structure fails when multiple EIFs from same software timeout in parallel which leads to the crash. [PR776497](#)
- Maximum number of supported IPsec tunnels might depend on networking activity as well. Under heavy networking activities, while DPD (Dead Peer Detection) is enabled, the maximum number of supported IPsec tunnels can drop to about 1800. [PR780813](#)
- If you set aggregated Ethernet interface ipfix sampling, IPv6 egress flow samples get snmp index of member link. Flow samples of IPv4 ingress, egress and IPv6 ingress do not experience this problem. They get snmp index of the aggregated Ethernet interface. [PR791619](#)
- In the Adaptive Service PIC (Service PIC II) scenario, configure the command "root@user# set services service-set <service-set-name> stateful-firewall-rules", because the command is not supported by 12.1R4, so Adaptive Service PIC goes offline. [PR819833](#)
- When rollback from v9 to v5 is done, Sampling logic was not rolling back, as sampling registers are not getting released from Packet Forwarding Engine and because in v5 the sampling is Routing Engine based it was not working. [PR824769](#)
- When DHCP subscribers login and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: GENCFG: op 24 (Lawful Intercept) failed; err 5 (Invalid)." [PR837877](#)
- SIP call forwarding may fail when NAT is used between parties even though the SIP ALG is in use. [PR839629](#)
- If flow-tap or radius-flow-tap is configured and logging, dynamic flow control daemon (dfcd) may be leaking file descriptors. Over time these leaked file descriptors reach the limit and following error message will be seen. /kernel: kern.maxfiles limit exceeded by uid 0, please see tuning(7). Then routing protocol daemon (rpd) may crash and generate a core file. [PR842124](#)
- Junos OS release introduces the IKEv2 support and a stricter check on IKE/IPsec SAs proposal params. [PR843893](#)
- When DHCP subscribers log in and radius hands down flow-tap variables the following errors are seen in the log: "/kernel: rts\_gencfg\_dependency\_ifstate(): dependency type (2) is not supported." [PR864444](#)



- The Remote Circuit ID DTCP trigger (X-RM-Circuit-Id) is being enhanced to have support for embedded whitespace (\040). [PR867937](#)
- Any port or IP address value set in SIP VIA header for 'rport' and 'received' attributes will not be checked or translated by the SIP ALG. There is usually no impact from this to a voice call, the contact address inserted by the client in future requests will be the external one but this will not disrupt the SIP ALG. Some rare clients however may have some unexpected reaction that causes problem such as try to register 2 IP addresses, the internal one and the public one, in the same register message which is unsupported by the ALG and causes the message to be dropped. [PR869725](#)
- Any SIP MESSAGE request will be dropped by the SIP ALG, this type of request is unsupported from day one. This is rare type of request which will not prevent more usual SIP operations such as voice calls, but it may affect some instant messaging applications based on SIP. [PR881813](#)
- In the Session Initiation Protocol (SIP) Application Layer Gateway (ALG) with port block allocation enabled scenario ("user@root# set services nat pool <pool-name> secured-port-block-allocation block-size <block-size>"), a SIP call to be set up and the ports block are allocated for the media flows. When the SIP media flows time out, the APP mapping starts using another port block. But if not enough port block is allocated, the services Physical Interface Card (PIC) might crash. [PR915750](#)
- "replicate-services" config command-line interface (CLI) under "set services service-set ..." is a hidden command, but it can be seen according to "root@user# run show configuration services | display set". [PR930521](#)
- No SNMP trap generated when NAT or Flow sessions reach the threshold. [PR933513](#)
- On M/MX/T Series platform with service-Physical Interface Card (service-PIC) scenario, when the http redirect feature is enabled, due to TCP sessions are not timed out correctly, http redirect running out of memory after running for a while, and the packet will be lost. [PR933696](#)
- DNS multiple queries A and AAAA might cause the Service-PIC to restart. [PR943425](#)
- When sending traffic from Internet end, the software count is incorrect. [PR948583](#)
- Message type for if\_msg\_ifl\_channel\_delete should be lower severity and not an error. [PR965298](#)
- In the L2TP scenario with dual Routing Engines. After subscriber management infrastructure daemon (smid) being restarted, because the delete notification to backup Routing Engine might be lost, the subscriber database (SDB) information does not synchronize between master Routing Engine and standby Routing Engine. After Routing Engine switchover is executed, the Layer 2 Tunneling Protocol daemon (jl2tpd) might crash, and new L2TP subscribers are unable to dial. [PR968947](#)
- If a PPPoE/PPP user disconnects in the access network without the LAC/LNS noticing it to tear down the connection (also the PPP keepalive hasn't detected yet), and a second PPP request comes from the same subscriber on the L2TP tunnel (same or different LAC/tunnel), then a second route is added to the table having the next hop "service to unknown". [PR981488](#)

- The cflow export would cease due to memory exhaustion when flow-monitoring is enabled using Adaptive Services II PIC due to memory leak condition. While in this condition, user would see increments in "Packet dropped (no memory)" as below:  
user@node> show services accounting errors Service Accounting interface: sp-3/0/0, Local interface index: 320 Service name: (default sampling) Interface state: Accounting Error information Packets dropped (no memory): 315805425, Packets dropped (not IP): 0. [PR982160](#)
- On MX240/480/960 Series router with MS-DPC with "deterministic-port-block-allocation block-size" configuration. In rare condition, when the "block-size" is set to a larger value (in this case, block-size=16128), the Services PIC might crash. [PR994107](#)

### ***Software Installation and Upgrade***

- Filesystem corruption might lead to Routing Engine boot up failure. This problem is observed when directory structure on hard disk (or SSD) is inconsistent. Such a failure should not result in boot up problem normally, but due to the software bug, the affected Junos OS releases mount /var filesystem incorrectly. The affected platforms are M/T/MX/TX/TXP. [PR905214](#)

### ***Subscriber Access Management***

- The authdlib logout/terminate release notify request might experience a processing loop. [PR888281](#)
- If there are services active at the time of the subscriber logout, authentication service process (authd) will send a service-deactivate request to dynamic profile database (profile-db)., if for some reason the request fails, then authd will retry until successful. If profile-db continues to fail the request, the authd continues to try the cycles continue and the subscriber might be 'stuck' in terminating state forever. [PR925723](#)
- Configuration change of the IPv4 address range in address-assignment pool does not always take effect. [PR954793](#)

### ***User Interface and Configuration***

- Selecting the Monitor port for any port in the Chassis Viewer page takes the user to the common Port Monitoring page instead of the corresponding Monitoring page of the selected port. [PR446890](#)
- User needs to wait until the page is completely loaded before navigating away from the current page. [PR567756](#)
- The J-Web interface allows the creation of duplicate term names in the Configure > Security > Filters > IPV4 Firewall Filters page. But the duplicate entry is not shown in the grid. There is no functionality impact on the J-Web interface. [PR574525](#)
- Using the Internet Explorer 7 browser, while deleting a user from the Configure > System Properties > User Management > Users page on the J-Web interface, the system is not showing warning message, whereas in the Firefox browser error messages are shown. [PR595932](#)

- If you access the J-Web interface using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR603669](#)
- On the J-Web interface, next hop column in Monitor > Routing > Route Information displays only the interface address and the corresponding IP address is missing. The title of the first column displays "static route address" instead of "Destination Address." [PR684552](#)
- Protected sections of the group hierarchy do not have their protection status displayed correctly and are not prevented from adding new elements into existing groups. [PR717527](#)
- "annotate" was not valid under firewall filter then hierarchy level and displayed "No valid completions", and lead to the configuration could not be committed under "edit private" mode. [edit] user@router# show | compare [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] + /\* Don't process the packet here; it's IPv6, not IPv4. + \* Accept it and have it be processed by the IPv6 ACL. \*/ accept; syntax error. user@router# commit full [edit firewall family inet filter LOOPBACK-OUTBOUND term allow-ipv6 then] 'accept' outgoing comment does not match patch. [PR812111](#)
- In an aggressive provisioning scenario using scripts or automated tools, we recommend that you do not use rollback immediately after a successful commit. [PR874677](#)
- On configure->clitools->point and click->system->advanced->deletion of saved core context on "No" option is not happening at J-Web. [PR888714](#)
- When PIM is enabled via apply-groups to one routing-instance whose instance-type is not defined (no-forwarding type is set), incorrect constraint check of PIM will cause routing protocol daemon (rpd) to crash upon any configuration change later. [PR915603](#)
- If a configuration file which contains groups related configuration is loaded by command "load replace", a "commit confirmed" operation might fail. When this issue occurs, the new configuration is committed even if you do not confirm it within the specified time limit. [PR925512](#)

### 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. [PR32763](#)
- BGP community 0xFF04 (65284) is a well known community (NOPEER), but it is incorrectly displayed as "mvpn-mcast-rpt" in the cli command "show route". This is a show command issue only. No operational mis-behavior will be observed on the router/network. [PR479156](#)
- If a logical interface is taken out of VPLS or L2VPN Pseudowire Routing Instance and placed in protocol l2circuit, after the above configuration changes are done in one commit, routing protocol daemon (rpd) crashes and dumps core. [PR872631](#)

- When a receiver already receiving multicast traffic for a group leaves the group, router connected to the receiver sends a Prune upstream and starts its upstream Prune timer. When the egress PE receives the Prune it will withdraw Type-4 route. During this time, if we 'clear pim join instance vrf' or (set routing-instances vrf protocols pim disable/enable) is done on egress PE and when the Receiver joins the group again, egress PE receives PIM Graft message but, drops it because it does not have matching SG state. This results in egress PE not being able to get trigger to send Type-4 and thereby is not able to pull traffic from ingress. [PR888901](#)
- In the Next-Generation Multicast VPN (NG-MVPN) scenario, there are two timing scenarios which could lead to the P2MP sub-LSP not being formed properly. Firstly, the issue happens if a Type 4 route from a neighbor is received before the Type 1 route from that neighbor. This is a transient condition, in case of unicast route flaps there is a possibility that the ordering of the routes gets altered and a Type 4 is received before a Type 1. The handling of such situation is missed for RSVP signaled P2MP tunnel. Secondly, before type 1 neighbor route is received, a Type 4 route is received on the ingress PE and the corresponding Type 3 route does not exist. When create a Type 3 route, the Type 1 neighbor route is still not received, in this case, this Type 4 do not get added as leaf and hence the tunnel is not formed when Type 1 neighbor route is received later. [PR913685](#)
- The issue happens when the virtual routing forwarding (vrf) is configured "no-vrf-propagate-ttl" and the vrf import policy changes the local preference of the vrf route. With "no-vrf-propagate-ttl", BGP will resolve the primary l3vpn route and the vrf secondary route separately. The root cause is overwriting the route parameters of the second vrf route with the route parameters of the primary route. So when changes the local preference of the vrf route might not work. [PR935574](#)

### ***Resolved Issues***

#### ***Class of Service (CoS)***

- On MX Series router with non-Q DPC (in this case, DPCE 40x 1GE R), when the "interface-set" is configured on a non-Q DPC, then execute the command "show interfaces interface-set queue <interface-set-name>", the DPC might crash. [PR979668](#)

#### ***Forwarding and Sampling***

- VPLS mac-table doesn't get populated with mac of previous lt interface after replacing the lt interface in the configuration, which might cause CE connected to the lt interface to get isolated. [PR955314: This issue has been resolved.](#)
- When port-mirroring or sampling is configured, if a lot of route updates are happening in the system, the routing protocol convergence time might be long and packet loss might be observed. [PR963060: This issue has been resolved.](#)

### General Routing

- BFD packets sent from FPC (distributed mode) over normal physical interfaces are set with ttl 0 so that it gets decremented by 1 and becomes 255 once it is sent out on the wire. This behavior is not the case when the BFD packets are sent over IPsec routed tunnels where the packets are sent from the corresponding service PIC. In this case, the ttl should be set to 255 as no such decrement action takes place when it is sent from a service PIC. But in the current scenario, the ttl is set to 0 as a result of which the service pic drops the outgoing packet. This was an untested scenario till date.  
[PR808545: This issue has been resolved.](#)
- RPD crashed on backup Routing Engine when trying to compare gateways of two different types of nexthops, like table next hop which is installed in kernel for one route, router next hop which is selected in backup RPD. [PR828797: This issue has been resolved.](#)
- In a scenario with scale Routing Instances (RIs) configured, after deactivating/activating two RIs, routing protocol daemon (rpd) might try to free a specific pointer pointing to an incorrect structure that is actively in use. Then rpd process crashes and dumps core files. [PR870683: This issue has been resolved.](#)
- On MX Series routers containing multiple Packet Forwarding Engines such as MX240/MX480/MX960/MX2010/MX2020, with either MPC3E or MPC4E cards (MPC3 Type 3 3D/MPC4E 3D 2CGE+8XGE/MPC4E 3D 32XGE), if multicast traffic or Layer 2 flood traffic enters the router via these MPC3E or MPC4E line cards, these line cards may exhibit a lockup, and one or more of their Packet Forwarding Engines corrupt traffic towards the router fabric. [PR931755: This issue has been resolved.](#)
- In the high scale P2MP LSP environment, heap memory leak might occur when the LSP flaps. Then some P2MP LSPs might be not installed, so the traffic will lose.  
[PR979211: This issue has been resolved.](#)
- scale-subscriber "License Used" filed shows wrong value after GRES. [PR980399: This issue has been resolved.](#)
- OpenSSL library in Junos OS was patched to resolve CVE-2010-5298. [PR984416: This issue has been resolved.](#)
- On M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120 and MX Series with DPC. In a race condition, the Dense Port Concentrator (DPC) may crash when ifls get added to an ifl-set while that same ifl-set get deactivated/deleted in class-of-service. For example: # set interfaces interface-set interface\_set\_JTAC\_ge-3/0/0 interface ge-3/0/0 unit 100 # deactivate class-of-service interfaces interface-set interface\_set\_JTAC\_ge-3/0/0 # commit or (quick commit of following changes) # set interfaces interface-set interface\_set\_JTAC\_ge-3/0/0 interface ge-3/0/0 # commit # deactivate class-of-service interfaces interface-set interface\_set\_JTAC\_ge-3/0/0 # commit. [PR985974: This issue has been resolved.](#)

### ***Infrastructure***

- In multicast scenario with composite next-hop (NH), after sending IGMP join for multicast groups and starting multicast traffic from other device, in some conditions, multicast NHs are changed (i.e. not the old NHs), Packet Forwarding Engine could not find the NHs, then FPC might crash. [PR822061: This issue has been resolved.](#)

### ***Interfaces and Chassis***

- When the GE port is configured with WAN PHY mode, a "Zero length TLV" message might be reported from the port. This is a cosmetic issue. [PR673937: This issue has been resolved.](#)
- When using vrrp inheritance, the vrrp might get stuck in bring-up state if it uses wrong vrrp parent, which will affect all new added vrrp even with a correct parent. As immediate recovery, it is suggested to correct or remove these two interfaces: xe-1/0/0.67 and xe-1/0/0.1066 and restart the vrrp process with the command: restart vrrp. [PR820298: This issue has been resolved.](#)
- When the remote device is using Address and Control Field Compression (ACFC) PPP compression, router will drop the received specific packet as it is not able to locate the PPP header. This causes L2TP sessions to not get established. [PR926919: This issue has been resolved.](#)
- In an MX Series router, multicast traffic may not be forwarded to the "Downstream Neighbors" as reported by the command "show pim join extensive". There can be occasions where this traffic is blackholed and not forwarded as expected. Alternatively, there may be an occasion where multicast traffic is internally replicated infinitely, causing one or more of the "Downstream Neighbors" to receive multicast traffic at line rate. [PR944773: This issue has been resolved.](#)
- In the VRRP for IPv6 environment with "checksum-without-pseudoheader" configuration under protocols vrrp hierarchy. If no other configuration under protocols vrrp hierarchy, when the configuration "checksum-without-pseudoheader" is deleted, the operation doesn't take effect, the knob is not cleared from the Virtual Router Redundancy Protocol daemon (vrrpd). This issue might lead to the two routers in VRRP Master-Master state. [PR958924: This issue has been resolved.](#)
- Temperature top and bottom are swapped in show chassis environments output for Type3/Type4 FPCs of T-Series [PR975758: This issue has been resolved.](#)

### ***Platform and Infrastructure***

- When the transit traffic is hitting the router and the destination is a local segment IP which requires ARP resolution, it's mis-classified by the DDOS filter and an incorrect policer is applied. This leads to host queue congestion. [PR924807: This issue has been resolved.](#)
- On MX Series routers with DPC type FPCs running a 11.4 (or newer) Junos OS release disabling family inet with uRPF enabled on a logical interface might result in another logical interface on the router to drop all incoming IPv4 packets. The lookup index is calculated by taking the lower 16 bits of the logical interface index (also called the IFL index). In other words lookup index = IFL index MOD 65536. It is normal, valid and

expected to have logical interfaces which share the same lookup index. The problem described in this PR is not the fact that the lookup indexes are the same. Here is an example of two different logical interfaces on two different FPCs which share the same lookup index: Interface ge-1/1/0.0 has an IFL index of 141073 and a lookup index 10001: > show interfaces ge-1/1/0.0 Logical interface ge-1/1/0.0 (Index 141073) (SNMP ifIndex 2318) ^^^^^^ Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Input packets : 0 Output packets: 0 Protocol inet, MTU: 993 ^^^^ Flags: Sendbroadcast-pkt-to-re, uRPF ^^^^^ Addresses, Flags: Is-Preferred Is-Primary Destination: 1.1.1.0/30, Local: 1.1.1.1, Broadcast: 1.1.1.3 Protocol multiservice, MTU: Unlimited Flags: Is-Primary And interface ge-2/0/7.1647 has an IFL index of 10001 and a lookup index of 10001: > show interfaces ge-2/0/7.1647 Logical interface ge-2/0/7.1647 (Index 10001) (SNMP ifIndex 20551) Flags: SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.1647 ] Encapsulation: ENET2 Input packets : 0 Output packets: 0 Protocol inet, MTU: 8978 Flags: Sendbroadcast-pkt-to-re, uRPF, uRPF-loose Protocol multiservice, MTU: Unlimited In the example above if family inet is disabled on ge-2/0/7.1647 then ge-1/1/0.0 will start dropping all incoming packets silently. [PR936249: This issue has been resolved.](#)

- On M7i/M10i/M120/M320 platform, when performing a software upgrade the upgrade may fail at the verification stage and a 'checkpic' coredump generated. [PR946582: This issue has been resolved.](#)
- Current display of "cli> request chassis routing-engine hard-disk-test show-status" command for Unigen SSD identified by "UGB94BPHxxxxx-KCI" is incorrect and can be misleading when use for trouble shooting. For example, attribute 199 is display as "UDMA CRC Error Count" is actually "Total Count of Write Sector". [PR951277: This issue has been resolved.](#)
- Unable to modify dynamic configuration database after first commit. [PR959450: This issue has been resolved.](#)
- Certain combinations of Junos OS CLI commands and arguments have been found to be exploitable in a way that can allow root access to the operating system. This may allow any user with permissions to run these CLI commands the ability to achieve elevated privileges and gain complete control of the device. Refer to JSA10634 for more information. [PR965758: This issue has been resolved.](#)
- If a router has DDoS protection enabled, when DDoS attack happened, the router might not be processing a DDoS event properly, which results in kernel crash. [PR987193: This issue has been resolved.](#)
- The non-first IP fragments containing UDP payload may be mistakenly interpreted as PTP packets if the following conditions are met: - the byte at the offset 9 in the IP packet contains 0x11 (decimal 17) - UDP payload - the two bytes at the offset 22 in the IP packet contain the value 0x01 0x3f (decimal 319; byte 22=0x01 and byte 23=0x3f) - PTP protocol The mis-identification of the packet as PTP will trigger the corruption of the fragment payload. [PR1006718: This issue has been resolved.](#)

### ***Routing Protocols***

- RPD process might core due to PIM join/prune internal processing. [PR817623: This issue has been resolved.](#)
- The routing protocol daemon (rpd) might generate a core file when multiple BGP sessions to neighbors in the same BGP peer group are forced to close status. Possible triggers are network reconvergence events causing the BGP sessions to go down, activating and deactivating the BGP protocol or the BGP peer-group configuration. [PR823346: This issue has been resolved.](#)
- In inter-AS option-C Layer 3 MPLS VPN scenarios, routing protocol process (rpd) might crash when multihop EBGP is configured between the 2 autonomous system boundary routers (ASBRs) and static LSPs are configured between the ASBRs to resolve the indirect nexthops. [PR869488: This issue has been resolved.](#)
- Forwarding cache limit is not properly meeting threshold after configuration change when configured per address family. [PR980578: This issue has been resolved.](#)

### ***Services Applications***

- Clearing the stateful firewall subscriber analysis causes the active subscriber count to display a very huge number. The large number is seen because when a subscriber times out the number of active subscribers is decremented. If it is set to zero using the clear command, then a decrement would give an incorrect result. There is no impact to the overall functionality and the fix will be present in 14.1R2. [PR939832: This issue has been resolved.](#)
- In H323 ALG with CGNAT scenario, the MS-PIC might crash when the ALG is deleting an H323 conversation due to the deleting port is outside of allocated NAT port-block range. [PR982780: This issue has been resolved.](#)
- On M/MX/T Series routers (platforms) with Services PIC with dynamic-nat44 translation-type configured, when the flows are cleared the IP addresses in use are never freed. This issue is present in Junos OS Release 11.4R7 and all more recent releases without this fix. [PR986974: This issue has been resolved.](#)



---

## Previous Releases

---

### Release 12.2R8

#### Forwarding and Sampling

- When MAC addresses move, Layer 2 address learning process (l2ald) will be called and produces some other child processes, the child processes cannot be terminated. Then maximum process limitation is hit and the Routing Engine is locked up. [PR943026: This issue has been resolved.](#)

#### General Routing

- Planes might go into faulty state during the SCB initialization when the SERDES on the SF chip failed to come up. [PR839509: This issue has been resolved.](#)
- MXVC /kernel: rts\_ifstate\_client\_open: Number of ifstate clients have reached threshold, current = 63 maximum = 63. [PR894974: This issue has been resolved.](#)
- This PR addresses a timing issue, which happens when "no-vrf-propagate-ttl" is configured in the routing-instance config. When this configuration is present, it might sometimes create a situation where the route selection happens of a route which is yet to be resolved in secondary vrf table, which results into a RPD core. [PR917536: This issue has been resolved.](#)
- MX80 routers now support CLI command "show system resource-monitor summary". [PR925794: This issue has been resolved.](#)
- When an SNMP walk is performed to query the native VLAN (mib-2.17.1.4.5.1...: dot1qPvid) or the logical type (trunk or access) of the interface (mib-2.17.1.4.3.1.5...: dot1qPortVlan), the SNMP walk might cause a memory leak on the Layer 2 address learning process (l2ald), and the process might crash with a core file generated. [PR935981: This issue has been resolved.](#)
- Master Routing Engine reboot due to "panic: pfe\_free\_peer: not in peer proxy process context" Trigger: replacement of backup Routing Engine. [PR936978: This issue has been resolved.](#)
- With scaled configuration of ATM VCs (~ 4000 VCs) on a single MIC-3D-8OC3-2OC12-ATM ATM MIC, the MIC might crash. The crash is not seen with lower scale (i.e. less than 3500 VCs per MIC). [PR947434: This issue has been resolved.](#)

#### Interfaces and Chassis

- DCD reports error when configuring hierarchical-scheduler on MX80 with QX chipset. This is cosmetic error and it should not have functional impact. [PR807345: This issue has been resolved.](#)
- Incorrect Detection timestamp in "show chassis fabric reachability" [PR811846: This issue has been resolved.](#)
- This is a timing issue in Multi-Chassis Link Aggregation (MC-LAG) active-standby scenario, when sequence of synchronize events executed on both routers with some delay during frequent MC-AE flaps with downstream device then MC-LAG in both

routers remain in active-active/standby-standby state. [PR885671: This issue has been resolved.](#)

- Flapping MLPPPoLNS (multiple ppp over L2TP network server) subscribers might cause logical interface (ifl) index leak, which results in subscribers being unable to connect or very slow to connect to MX Series router. [PR886474: This issue has been resolved.](#)
- In MX Series multichassis link aggregation group (MC-LAG) scenario, if change the configuration from MC-LAG to non-MC-LAG and then to MC-LAG, the Layer 2 address learning daemon (l2ald) might crash due to the incorrect media access control (MAC) address processing. [PR893842: This issue has been resolved.](#)
- In Point-to-Point Protocol over Ethernet (PPPoE) scenario, if some PPPoE session was added and deleted, after performing Routing Engine switchover operation, the Broadband Remote Access Server (BRAS) might fail to allocate PPPoE session IDs on interFace Descriptor (ifd). [PR896946: This issue has been resolved.](#)
- When you make changes on the Virtual Router Redundancy Protocol (VRRP) enabled interface who is master, such as disable interface and re-enable it/deactivate the interface and re-activate it/interface state down and up, or when you restart VRRP process, Address Resolution Protocol (ARP) request packet might be sent out with incorrect VRRP MAC address (00:00:5e:00:01:00) over the interface. Issue is because of VRRP process not configuring VRRP group id and state when it is in transition state (in the transition state from backup to master of the whole "Idle->backup->master" process). This issue is not specific to MX Series platform. But as this is a timing issue, it was more frequently seen on MX Virtual Chassis (MXVC) scenario. As workaround, you can disable skew timer at the [edit protocols vrrp] hierarchy level by "lab@R0#skew-timer-disable" if there is only one master router and one backup router in the network deployment (statement introduced in 12.2 and not support on every platform) or use virtual IP address same as interface address. [PR908795: This issue has been resolved.](#)
- "Too many I2C Failures" alarm happens when a FRU (in this case: PWR-MX960-4100-AC-S) experienced 6 consecutive i2c read/write failures. While the PEM is still providing power to the chassis, chassisd daemon cannot read/write information from the PEM until it is reseated. Some enhancements have been made for this MX960 HC AC PEM: 1. PEM i2c bus hang avoidance 2. Junos OS recovery from a hung i2c bus 3. noise reduction This Junos OS eliminates the need for the PEM FW upgrade, and at the same time is 100% compatible with those PEMs which have been upgraded. [PR928861: This issue has been resolved.](#)
- PCS statistics counter(Bit errors/Errored blocks) not working on Mammoth PIC(xge). [PR942719: This issue has been resolved.](#)

### Layer 2 Features

- In a protocol-mastership transition, the ksyncd process might fail to clean up the kernel VPLS routing tables due to dependencies such as VLANs not being cleaned up first, leaving the tables in an inconsistent state. ===== BACKGROUND ===== A global GRES, which will cause a master Routing Engine to transition to backup, WILL require all Kernel state to be cleaned so that it can start a fresh resync from the new master. Ksyncd is tasked with cleaning up Kernel state. On cleaning routing tables, if any table has a non-zero reference count, it will return "Device Busy" to the ksyncd. Ksyncd will try 5 successive cleanup attempts after which it will trigger a live Kernel core. ===== PROBLEM ===== In ksyncd's kernel cleanup, the Bridge Domain mapped to a VPLS routing table is deleted AFTER an attempt is made to delete the route table. This is a catch-22 since BDs hold reference counts to the routing table. ===== FIX ===== Cleanup of VPLS routing tables should proceed bottom up in the following order: NextHop Deletes, User Route Deletes, Interface Deletes(ifd,ifl,iff), STP Deletes, Bridge Domain Deletes, Mesh Group Deletes and finally Routing Table delete. This ensures that when we get to routing table delete, all dependencies, that could hold a ref cnt to the routing table, are now gone. [PR927214: This issue has been resolved.](#)

### MPLS

- When static LSPs are configured on a node, RPD could assert upon committing a MPLS-related configuration change. Example: router> show system rollback compare 9 8 [edit protocols mpls] interface ae11.0 { ... } + interface as3.0 { + admin-group red; + } [edit protocols isis interface as3.0 level 2] ! inactive: metric 2610; The following error is seen in /var/log/messages in-relation to a static lsp, immediately following the above-mentioned configuration change: rpd[1583]: UI\_CONFIGURATION\_ERROR: Process: rpd, path: [edit groups STATELESS\_ARIADNE protocols mpls static-label-switched-path static-lsp], statement: transit 1033465, static-lsp: incoming-label 1033465 has already been configured by this or other static applications. [PR930058: This issue has been resolved.](#)
- In certain circumstance, the Junos OS rpd route flash job and LDP connection job are always running starving other work such as stale route deletion. These jobs are running as LDP is continuously sending label map and label withdraw messages for some of the prefixes under ldp egress policy. This is due to LDP processing a BGP route from inet.3 for which it has a ingress tunnel (the same prefix is also learnt via IGP) creating a circular dependency as BGP routes can themselves be resolved over a LDP route. [PR945234: This issue has been resolved.](#)
- When Packet Forwarding Engine fast reroute (FRR) applications are in use (such as MPLS facility backup, fast-reroute, loop free alternates), a flap of the primary path could be triggered due to an interface flap or by Bidirectional forwarding detection (BFD) session flap. However, this interface/session flap might lead to a permanent use of the backup path, which means the original primary path could not be active again. [PR955231: This issue has been resolved.](#)

### Network Management and Monitoring

- When there are mix of OIDs in a PDU Get request, the subtree information of certain OIDs can be NULL if the respective subagent hasn't registered completely the

corresponding MIB objects or if there is some other underlying issue which would cause the subtree information to be NULL. Accessing NULL information caused SNMP to crash. [PR779346: This issue has been resolved.](#)

- When you perform the below MIB Walk on interfaces, for some interfaces the ifLastChange value will show a value of zero. show snmp mib get ifLastChange. <SNMP ifIndex> will show a value of zero. ifLastChange. <SNMP ifIndex>=0. [PR886624: This issue has been resolved.](#)

### ***Platform and Infrastructure***

- In a MX-VC environment, in certain situations the inter-chassis traffic may not be equally balanced across all available Virtual Chassis Port (VCP) links after adding extra links. [PR915383: This issue has been resolved.](#)
- Upon the deletion of a routing-instance and subsequent commit, error logs are generated from each Type 1 - 3(non E3) based FPC. These logs are cosmetic and can be ignored. [PR964326: This issue has been resolved.](#)

### ***Routing Policy and Firewall Filters***

- Policy with Install-nexthop lsp may not work as expected when there is an LSP path change triggering route resolution. [PR931741: This issue has been resolved.](#)

### ***Routing Protocols***

- With BGP import policy as next-hop peer-address, if the local router receives inet (or inet-vpn) flow network-layer reachability information (NLRI), routing protocol process (rpd) might crash. Junos OS is designed to create a fictitious nexthop for inet flow and inet-vpn flow families as they don't send/expect-to-receive nexthops. So in this case when the import-policy set a non-null next-hop for the received inet (or inet-vpn) flow route, it could not handle properly which might result in rpd crash. [PR966130: This issue has been resolved.](#)

### ***Services Applications***

- During a rare scenario, switchover on another sp interface can crash a service PIC when running a traffic in hairpinning scenario. [PR945114: This issue has been resolved.](#)

### ***Release 12.2R7***

#### ***Class of Service (CoS)***

- After swapping MPC2E-3D-Q card with MPC2E-3D-EQ card, an interface is still running out of queues with only 32k queues in use. [PR940099: This issue has been resolved.](#)

#### ***Forwarding and Sampling***

- When pfd gets restarted during a period when pfd is communicating with mib2d, because the communication sockets have been terminated and failed to be re-opened after pfd came back up again, mib2d might crash and generate a core file. The core files could be seen by executing CLI command **show system core-dumps**. [PR919773: This issue has been resolved.](#)

### ***General Routing***

- Planes might go into faulty state during the SCB initialization when the SERDES on the SF chip failed to come up. [PR839509: This issue has been resolved.](#)
- Ipv6 address syntax on rpd log is violated of RFC 5952. For example, 2002:db8:0:0:1:0:0:1 must be logged as 2002:db8::1:0:0:1 in the logs, but it's logged as 2002:db8:0:0:1::1. 2001:0:0:0:db8:0:0:1 must be logged as 2002::db8:0:0:1 in the logs, but it's logged as 2001:0:0:0:db8::1. The fix is available in 11.4R10, 12.1R9, 12.2R7. [PR840012: This issue has been resolved.](#)
- In subscriber management environment with auto-sensed VLAN configured, in a rare case, after some configuration changes made, kernel crash is observed leading to Routing Engine reboot. The issue is identified as an interface which is not initialized properly getting packets. [PR878921: This issue has been resolved.](#)
- Following a global GRES event, the new Master(VC-Mm) will expect relayd to reconnect to it in less than 40 seconds. However under high scale, such as with 54k dual-stack(v4v6) or 110k+ single-stack DHCP subscribers, owing either to a slow

relayd(relay daemon) control connection to the Kernel, or due to slow pfe reconnects to relayd, we are not able to meet the 40 seconds timer requirement causing subsequent FPC reboots and traffic loss. [PR891814: This issue has been resolved.](#)

- When GRES and ARP purging is enabled, frequent route flapping, route entry and nexthop fail to sync up between master Routing Engine and backup Routing Engine. So when master Routing Engine would like to add a new nexthop but sees backup Routing Engine has already found a nexthop with same destination, it makes backup Routing Engine reboot and crash on both Routing Engines. [PR899468: This issue has been resolved.](#)
- 100G Ethernet interface (Finisar FTLC1181RDN3-J3) on T4000 type-5 FPC may flap once after bringup . The solution is changing the register bandwidth. [PR901348: This issue has been resolved.](#)
- RPD on backup Routing Engine might hit out of memory condition and crash if BGP protocol experiences many flaps [PR904721: This issue has been resolved.](#)
- When adding the "no-tunnel-services" knob under VPLS protocols of routing-instances, during the processing gap of the new knob, if routing protocol process (rpd) restarts (i.e rpd crashes), logical interfaces with VPLS family do not show up, and there are no logical interfaces available for the corresponding VPLS routing instances. Hence VPLS connections might be down (stuck in LD state) and can not be recovered automatically. [PR912258: This issue has been resolved.](#)
- In multi-router Automatic Protection Switching (APS) scenario, the laser of the protection link might be turned off and never come back on when the ATM (at-) interface of the Circuit Emulation MIC flap or the MIC restart. In such conditions, if the working link goes down, APS fails to switch traffic to the protection link. [PR917117: This issue has been resolved.](#)
- With out this PR fix, MS PIC running E Junos OS will not be able to sync timezone configuration with Routing Engine. [PR926488: This issue has been resolved.](#)
- The output of "show subscribers summary slot" is incorrect. [PR926508: This issue has been resolved.](#)
- tcp\_inpcb buffer leak in ADC and TLB service pics [PR934768: This issue has been resolved.](#)
- LNS drops the LCP Compression Control Protocol (CCP) packet silently comes from L2TP tunnel. [PR940784: This issue has been resolved.](#)

#### ***High Availability (HA) and Resiliency***

- With minimal flow configuration, if graceful Routing Engine switchover (GRES) is not enabled, routing protocol process (rpd) crashes during shutting down the rpd process due to missing safety checks. The core files could be seen by executing CLI command "show system core-dumps". [PR852766: This issue has been resolved.](#)
- During every failover of redundancy-group 0, the /etc/ssh and /var/db/certs directories are copied from primary node to secondary node. However, the directories are not copied correctly and nested directories such as /etc/ssh/ssh, /etc/ssh/ssh/ssh are created. [PR878436: This issue has been resolved.](#)

**Infrastructure**

- ksyncd disconnects after resync during unified ISSU. [PR882027: This issue has been resolved.](#)

**Interfaces and Chassis**

- An MX Series router may cosmetically log "Bottom Fan Tray Unable to Synch". [PR833047: This issue has been resolved.](#)
- Tx and Rx Spanning-tree BPDU stopped intermittently during unified ISSU. [PR849201: This issue has been resolved.](#)
- Reboot after panic: xe-0/1/0: bitstring index 7 not empty for 01:00:5e:00:00:01 ( fix needed for MPC/MIC) [PR905417: This issue has been resolved.](#)
- In multicast over AE scenario, if there is a different order of child IFLs (logical interface) under parent AE at master Routing Engine and backup Routing Engine, then after Routing Engine switchover, multicast traffic might get lost. [PR915440: This issue has been resolved.](#)
- For IQ2 PIC, when the setting shaping rate is too high, when configured it with "set chassis fpc 0 pic 1 traffic-manager logical-interface-base-shaping-rate 16" and this will reset the shaping rate to 1Gbps. The corresponding messages are logged in debug level. In the fix, it is corrected into info level. [PR920690: This issue has been resolved.](#)
- Unified ISSU fails on upgrade to 11.4R5.7. with the following message Logged messages: MIC 4/0 will be offlined (In-Service-Upgrade not supported) MIC 4/1 will be offlined (In-Service-Upgrade not supported) Do you want to continue with these actions being taken ? [yes,no] (no) yes error: /usr/sbin/indb failed, status 0x200 error: ISSU Aborted! Chassis ISSU Aborted ISSU: IDLE Issue happens when a MIC-3D-4OC3OC12-IOC48 card is offline via CLI and removed from the chassis prior to the ISSU. [PR923569: This issue has been resolved.](#)
- Traffic which uses MPLS next-hops enters bridge-domain via IRB interface and if forwarding next-hop moves from non-aggregate interface to aggregate interface (MAC move), the MPLS next-hops are not correctly programmed in the Packet Forwarding Engine. The child next-hop of the aggregate interfaces are missing. Once IRB MPLS next-hop moves from aggregate interface to non-aggregate interfaces are not affected. IPV4 traffic is not affected. [PR924015: This issue has been resolved.](#)
- In PPPoE subscriber management environment, when PPP daemon is receiving an LCP packet with an invalid code ID and without any option, jpppd process crashes with a core file generated. [PR929270: This issue has been resolved.](#)
- After APS switchover, duplicate packets might be received from the backup circuit under SONET APS configuration with channelized enhanced intelligent queuing (IQE) interface. [PR930535: This issue has been resolved.](#)

### **Layer 2 Features**

- "show snmp mib walk ascii jnxVpnIfStatus" doesn't work for BGP VPLS when there is incompleted BGP VPLS instance configuration or LDP VPLS instance. [PR918174: This issue has been resolved.](#)

### **Layer 2 Ethernet Services**

- In MX Virtual Chassis (MXVC) scenario, under high scale system environment (many Aggregated Ethernet interfaces, many logical interfaces), after performing global graceful Routing Engine switchover (GRES) by CLI command "request virtual-chassis routing-engine master switch", the Link Aggregation Control Protocol (LACP) state of access Link Aggregation Group (LAG) interface might change and therefore resulting in traffic loss. [PR885013: This issue has been resolved.](#)

### **MPLS**

- The RPD process may crash when executing the command "clear mpls lsp name <lspname>" or "monitor label-switched-path <lspname>". [PR756551: This issue has been resolved.](#)
- IPv6 traceroute may not show some hops for scenarios where 1) Two LSPs are involved. 2) INET6 Shortcuts are enabled. In such scenarios, hops that are egress for one LSP and ingress for the next LSP in the traceroute do not show up. This was a software issue with icmp error handling for packets with ipv6 payload having a ttl of 1. [PR899283: This issue has been resolved.](#)

### **Platform and Infrastructure**

- With l3vpn composite next-hops configured and 3 or more odd number of core uplinks every l3vpn route deletion will syslog the following error messages. [LOG: Err] JTREE: (jt\_mem\_free) size 0 for addr 1595452, seg 1, inst 0 [LOG: Emergency] Multiple Free :jt\_mem\_free There is no operational impact. An even number of core-uplinks will not trigger such error logs. [PR786993: This issue has been resolved.](#)
- The FPC contains "1x CHSTM1 IQ, SMIR" PIC may crash when the "1x CHSTM1 IQ, SMIR" PIC online/offline. The core files could be seen by executing CLI command "show system core-dumps". The root cause of the problem is the bottom driver defect: when the PIC online/offline, the SONET/SDH framer in this PIC will stick in reading the J0 byte received and hogging for more than 2500 milliseconds, then leads to the crash. [PR806317: This issue has been resolved.](#)
- Packet Forwarding Engine might crash when receiving TCP packets with wrong IP header (IP packet length is shorter than IP header length). With the fix, Junos OS discards the invalid packet as layer-3 incomplete on the interface level and avoids the crash. [PR817318: This issue has been resolved.](#)
- The system MAC address is not getting saved in a unified in-service software upgrade (unified ISSU) blob and it is not getting programmed again by the Routing Engine when the Packet Forwarding Engine re-connects. The hash seed is generated by using the system MAC address and since it is not saved in a unified ISSU blob, after a unified ISSU it is 0 and the hash seed is generated using that. If a FPC reboot, then it will get the correct system MAC address and generate the hash seed based on that. This will



cause different FPCs in the system to have different hash seeds and could cause AE multicast traffic loss if the ingress and egress FPCs have different hash seeds. [PR915933: This issue has been resolved.](#)

- In subscriber management scenario, memory leak might occur when the firewall fast-update-filter feature is configured, and it will impact any new subscriber login. Such memory leak can be seen with following command, `root@router> show chassis fpc Temp CPU Utilization (%) Memory Utilization (%) Slot State (C) Total Interrupt DRAM (MB) Heap Buffer 0 Online Absent 8 0 1024 70 << 13 1 Online Absent 8 0 1024 29 13` [PR926808: This issue has been resolved.](#)
- On an MX Series chassis running in "enhanced-ip" mode, multicast traffic forwarded on MX Series FPCs might be corrupted if PTP protocol packets (Precision Time Protocol) are also forwarded on the same Packet Forwarding Engine. [PR932471: This issue has been resolved.](#)
- "Total errors" counter of MAC statistics on MX DPC(ge/xge) is always 0. [PR942183: This issue has been resolved.](#)

### ***Routing Policy and Firewall Filters***

- Junos OS releases with a fix for PR/706064 have a regression where the vrf-import policy sanitation logic is faulty. A "# commit check" will fail when the first term references a 'target' community and the second term references an 'origin' community. This should pass the check. [PR911350: This issue has been resolved.](#)

### ***Routing Protocols***

- When the IPv6 address on fpx0 is active during boot up, the joining of the all-router group causes the kernel to create a ff02::2 route with a private nexthop, which is not pushed to the Packet Forwarding Engine. When a non-fpx0 interface is active later, the private nexthop will be shared by the non-fpx0 interface as well, resulting in packet drops destined to ff02::2 on the non-management interface. [PR824998: This issue has been resolved.](#)
- There is improper </route-family> tags added to all "multicast route summary" commands when we perform command such as `show multicast route summary | display xml`. [PR859104: This issue has been resolved.](#)
- Rpd may crash on the new master Routing Engine after Routing Engine switchover. The issue is NSR related, and it happens due to the bad BGP route data structure on the backup Routing Engine. [PR885305: This issue has been resolved.](#)
- BGP "accepted-prefix-limit" feature might not work as intended when it is configured together with "damping". Root cause of this issue is that when BGP module count the maximum routes accepted from BGP neighbor, it doesn't count the accepted BGP routes which in damping status. So when these damping routes are reused, the total number of received BGP routes exceeds the configured value for "accepted-prefix-limit". [PR897124: This issue has been resolved.](#)
- In multicast scenario with PIM enabled, when you configure both static RP mapping with override knob and dynamic RP mapping (such as auto-RP) in a single routing instance, allow the static mapping to take precedence for a given group range, and

allow dynamic RP mapping for all other groups, but a software defect cause that RP is selected based on dynamic RP mapping address, instead of accounting for this static override knob. [PR912920: This issue has been resolved.](#)

- DR sends a delayed ACK to the LSA on the interface on which the LSA is flooded. This leads to BDR sending only directed ACK to DR, DR-Other is therefore not receiving this ACK and is hence retransmitting the LSA to BDR. [PR914803: This issue has been resolved.](#)

### ***Services Applications***

- NAPT: Packet Forwarding Engine side reports port range start from 512 causing napt mib counter to be wrong. This fix makes the port range in pfe start from 1024. [PR828450: This issue has been resolved.](#)
- In Carrier Grade NAT scenario, MS-PIC might crash and generate a core file when Port Block Allocation (PBA) block size is relatively big (8192 ports per block). This issue usually happens when a new block needs to be allocated because the block currently is exhausted. [PR874500: This issue has been resolved.](#)
- In rare conditions with large number of traffic flows ( like NAT and IPsec flows ), the Service PIC may get stuck or crash as a result of prolonged flow-control assertions towards the Packet Forwarding Engine. In order to trigger this issue, many Compute CPUs inside the Service PIC should be overloaded. This will never happen under normal operation, where CPUs can handle large amounts of traffic without any issues. [PR900227: This issue has been resolved.](#)
- In a CGNAT environment, active FTP operations fail when there is latency issue in network. When TCP retransmission, FTP ALG is not translating any fields in the Request: PORT command. As a result server tries to establish the data flow to the private IP address and to a wrong TCP port and it fails as expected. [PR916376: This issue has been resolved.](#)
- In Carrier Grade Network Address Translation (CGNAT) with high memory utilization environment (Memory is in yellow zone and use CLI "show services service-sets memory-usage" to check), this crash might be seen in hairpinning scenario where Endpoint Independent Filtering (EIF) is enabled and the initial packet of a specific flow that hits the MS-DPC is dropped by an ALG due to various reasons (malformed or non complying packet/headers). [PR918663: This issue has been resolved.](#)
- In Carrier Grade NAT (CGNAT) environment, during heavy setup rate of CGNAT flows, inter-chassis stateful High Availability (HA) sync flaps and then keepalive messages are lost, as there is no control flow prioritization configured. HA sync connection keeps disconnecting. After a long period of time PIC silently reboots. Following syslog message might be seen when issue occurs: ROUTER-RE0 (FPC Slot 2, PIC Slot 0) PFEMAN: Lost contact with master routing engine PFEMAN: Forwarding will cease in 4 minutes, 59 seconds ROUTER-RE0 (FPC Slot 3, PIC Slot 1) PFEMAN: Lost contact with master routing engine PFEMAN: Forwarding will cease in 4 minutes, 59 seconds. [PR920723: This issue has been resolved.](#)
- In Carrier Grade Network Address Translation (CGNAT) environment whenever an inbound UDP packet did not hit any rule, a check is performed whether the destination ip and port match any SIP registration. If this check is successful and 'learn-sip-register' is enabled (which is the default in the junos-sip application), if packets are counted

as SIP ALG parsing errors, no flow is created and the packet will be forwarded without any transformation. In the case of NAT, the destination address will remain within the NAT pool and the packet will keep coming back to the service PIC, causing a routing loop and high CPU utilization. [PR923630: This issue has been resolved.](#)

- If multiple service sets with different number of NAT rules/pools are configured, Services PIC might crash when SNMP walk is performed on jnxSrcNatStatsTable. [PR928169: This issue has been resolved.](#)
- When tcp session is initiated from inside client and three way handshake is not completed due to the fact that client did not ack the syn-ack send from the server, service pic will send a tcp reset to the server after the timer expires. In this case tcp reset is sent on the wrong direction, instead of sending on the outbound direction to the server, service pic will send it in the inbound direction. This PR fixes this issue. No service impact is seen because of this. [PR931433: This issue has been resolved.](#)
- Interim-logging is now supported with NAT64 on microkernel (MS-DPC) platforms. The same pba-interim-logging-interval knob under 'service-options' under the service interface will enable the feature for NAT64 as well. [PR935606: This issue has been resolved.](#)
- When FireWall (FW) flows trying to create a new pair of flows, some resource was used by other FW flows, then service pic crashes. [PR940014: This issue has been resolved.](#)

### ***Software Installation and Upgrade***

- In this case, since the high level package (i.e. jinstall) is signed, the underlying component packages are not required to be signed explicitly. However the infra was written such a way to display warning message if the component package is not signed (i.e. jpfe). [PR932974: This issue has been resolved.](#)

### ***Subscriber Access Management***

- If there is secureid configuration present on the chassis, when the validate phase of "request system software add" runs, the netstat might crash due to system can not load the SecureID module during syntax checking. The generation of the core file has no effect on the verification results, and does not adversely affect the upgrade/downgrade operation. [PR911232: This issue has been resolved.](#)

### ***VPNs***

- Configuration version (child rpd) of rpd generates a core file when doing a commit or commit check. [PR930080: This issue has been resolved.](#)

### ***Release 12.2R6***

#### ***Class of Service (CoS)***

- During addition/deletion or just deletion of interfaces with configuration for shared scheduler, some portion of memory is not reclaimed back normally. So continuous addition/deletion of these interfaces results in memory depletion, packet loss and other issues. [PR890986: This issue has been resolved.](#)

#### ***Forwarding and Sampling***

- Outbound control traffic is not counted by accounting-profile which applied to logical interfaces of AE (Aggregated Ethernet). This is a variation of the PR562964. [PR866181: This issue has been resolved.](#)
- lab@T1600-2\_Critical\_VZB\_Manjit> show services accounting flow-detail destination-prefix 20.1.1.2/32 Service Accounting interface: sp-2/0/0, Local interface index: 147 Service name: (default sampling) Interface state: Accounting Protocol Input Source Source Output Destination Destination Packet Byte Time since last Packet count for Byte count for interface address port interface address port count count active timeout last active timeout last active timeout udp(17) xe-0/0/3.0 10.1.1.2 whois++(63) xe-0/0/2.0 20.1.1.2 whois++(63) 1075917 49492182 00:17:55 1780922 81922412 tcp(6) xe-0/0/3.0 10.1.1.2 0 xe-0/0/2.0 20.1.1.2 0 106479 4898034 00:01:46 1835070 84413220 [PR881629: This issue has been resolved.](#)
- In scaled MPLS scenario, when LSP path switchover happens, sample process deletes sampling parameters from the Packet Forwarding Engine and as a result of that Packet Forwarding Engine stops exporting flows to the collector. [PR891899: This issue has been resolved.](#)

### General Routing

- Only 94 GRE(plain) sessions are in Established state after chassisd restart. [PR801931: This issue has been resolved.](#)
- authd reports syntax error, although the syntax is correct, when trying to activate service profile for subscriber and fails to activate the service. [PR883065: This issue has been resolved.](#)
- After deactive or delete NSR configuration, the Routing Engine might become non-responsive due to the exhaustion of kernel buffer with following messages: /kernel: Mbuf: High Utilization Level: (Low) Throttling low priority requests (10 ms) /kernel: Mbuf: High Utilization Level: (Medium) Throttle low priority requests (150 ms) /kernel: Mbuf: High Utilization Level: (High) Block low priority requests You can get the kernel buffer usage by CLI command "show system buffers". [PR886083: This issue has been resolved.](#)
- In MX Series virtual chassis (MXVC) scenario, nexthop statistic requests such as "show mpls lsp statistics" from the Kernel to the Packet Forwarding Engines have to go via relay daemon. Under scaled configurations, the nexthop statistic requests message that is being sent to the Bm-Routing Engine is bigger than the max allowed size, causing kernel on Mm-Routing Engine to crash with core files generated, then Mm-RE goes down. [PR886864: This issue has been resolved.](#)
- When multiple framed-route(type-22) AVPs are present in Radius access accept message, the router will install only the first route into the routing table. [PR891036: This issue has been resolved.](#)
- When a bgp route is resolved using a next-hop that is also learned in bgp (i.e. there are multiple levels of next-hop resolution) and bgp multipath is also used, during a route churn next-hop for such a bgp route could be incorrectly programmed. [PR893543: This issue has been resolved.](#)
- In subscriber management environment, in a rare case, VLAN auto-sensing daemon (autoconfd) might crash and create a core file due to Session Database (SDB) is inaccessible. [PR899747: This issue has been resolved.](#)
- Some ATM interfaces may stay down after flapping the Circuit Emulation MIC. [PR900926: This issue has been resolved.](#)
- bootp configuration on TXP platform referencing routing-instance fails to commit. [PR906713: This issue has been resolved.](#)
- VCMm-power down creates stale vlan demux0 entries at the Packet Forwarding Engine level. [PR908027: This issue has been resolved.](#)

### ***High Availability (HA) and Resiliency***

- In certain systems configured with GRES, there is the possibility for the master and the backup Routing Engine to reach an inconsistent view of installed state. This fault may be exposed if the master Routing Engine experiences a mastership watchdog timeout at a time when it is not in sync with the backup Routing Engine for a particular piece of state. In practice, this possibility exists only for a short time period after an Routing Engine mastership change. Under such conditions, a replication failure may cause the backup Routing Engine to panic. If the failure is seen, the backup Routing Engine will recover on restart. In 11.4 and 12.1 releases without this fix, the fault may be experienced on any GRES-enabled, non-multichassis configuration on a T Series router. For 12.2 and later releases without this fix, the fault may be experienced on any GRES-enabled, non-multichassis configuration on a T Series or MX Series router. [PR910259: This issue has been resolved.](#)

### ***Infrastructure***

- When a sonet interface with PPP encapsulation is used as forwarding next hop for the IPv6 remote router loopback address on IPv6 BGP sessions, if the sonet link is down, the IPv6 BGP session might flap at same time although there is valid route via other interface. [PR863462: This issue has been resolved.](#)
- Kernel may crash when delete routing instance under the donor and unnumbered address borrower scenario. When the deleting for the donor is before the deleting of the corresponding unnumbered borrower, in this window, the donor interface does not have an address, arp processing over the borrower interface during this window may trigger the crash. The core files could be seen by executing CLI command "show system core-dumps". [PR880179: This issue has been resolved.](#)
- When multicast is running on a multi-chassis environment, during flapping of 224/4 or ff00/8 pointing to mResolve(NH), the LCC master might get replication error which causing all FPCs going offline. This flapping of resolve route for multicast can occur because of any of the following reasons: enabling or disabling multicast, deletion of resolve route, or routing restart. [PR897428: This issue has been resolved.](#)
- In a multihop IPv6 BGP session scenario, after configuring single-hop BFD session on the multihop IPv6 BGP neighbor, kernel might try to access a NULL pointer, causing kernel to crash and create a core file. [PR898153: This issue has been resolved.](#)
- Checksum error seen on ICMP reply when 'sequence, data' field in request set to '0'. [PR898487: This issue has been resolved.](#)

### ***Interfaces and Chassis***

- Traffic loss is seen, Multiple inbound and outbound IPsec tunnels are created for a single SA during tunnel renegotiation after the lifetime expiry. [PR827647: This issue has been resolved.](#)
- IQ2 core is seen after unified ISSU and traffic will be lost for a while (about 40s). The crash happens during processing of scheduler free message which comes just after unified ISSU complete on IQ2. Then the heap structure is invalid causing panic. The fix is moving the process to unified ISSU sync stage. [PR845257: This issue has been resolved.](#)

- In a scenario of PPP sessions over L2TP tunnels, on L2TP network server (LNS), if authentication is none or if authentication is enabled but radius does not return any Framed-IP-Address/Framed-Pool, jpppd process is not setting the IP address key of subscriber to "255.255.255.254" thereby resulting in address allocation failure in authd process. Then the L2TP tunnels can not be established, hence subscribers can not login. When issue happens, the following logs of authd process could be seen: client type jpppd client type REQUESTING: OldStyle 0 OldStyleFilled 0 hint null network null client pool name. [PR849191: This issue has been resolved.](#)
- "Dump-on-flow-control" knob might not work correctly for RSP interfaces configured in "warm-standby" mode. After an RSP switchover, either manual or following a crash, the "dump-on-flow-control" flag might get cleared from the MS-PIC. [PR867394: This issue has been resolved.](#)
- M7i Routing Engine Crashed with last reboot reason panic:page fault and kernel core, after commit. [PR868212: This issue has been resolved.](#)
- Chassisd core generated on initializing process on MX-VC. [PR870457: This issue has been resolved.](#)
- In subscriber management environment, with dynamic-profiles configured for subscribers, if the routing instance returned from radius is not configured on BRAS, dynamic-profile add fails and there are some places the memory not freed, causing device control daemon (dcd) memory leak. The memory usage of dcd process can be observed by following command: user@router> show system processes extensive | match dcd PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND  
7076 root 1 97 0 1047M 996M select 6:05 2.88% dcd [PR880235: This issue has been resolved.](#)
- VC-Boot loop when installing new local backup Routing-Engine. [PR881906: This issue has been resolved.](#)
- While a duplicate interface address (IFA) is configured for two interfaces, software will accept that and create a error message like this:  
%CONFLICT-4-DCD\_PARSE\_WARN\_INCOMPATIBLE\_CFG: [edit interfaces ge-0/0/0 unit 0 family inet address x.x.x.x/xx] : Incompatible configuration detected : identical local address is found on different interfaces But at kernel side cannot accept duplicate IFA, and needs to delete the next-hop created for this operation. Due to code problem, the clean up doesn't remove the duplicated IFA under heavy kernel workload. And it will crash while trying to update this duplicated IFA to the Packet Forwarding Engine side. [PR891672: This issue has been resolved.](#)
- In dynamic PPPoE subscriber management environment, when MS-DPC card is added and "adaptive-services service-package laryer-2" is configured, while PPPoE subscribers log in, kernel might encounter a memory corruption, causing kernel to crash and generate a core file. [PR894440: This issue has been resolved.](#)
- On MX Series routers with MICs/MPCs, when PIC is configured with traffic-manager mode ingress-and-egress, after PIC offline, PIC detach does not clean up the corresponding entries completely. Subsequent PIC online results in corresponding entries add failure since previous entries are still intact, resulting in interface attach failure at the Packet Forwarding Engine level. Due to interface add failure, protocols on the interface never come up. [PR895305: This issue has been resolved.](#)

### Layer 2 Features

- Frames containing PPPoE encapsulated packets may get dropped when they need to be flooded and forwarded across the fabric in a bridged environment on MX Series routers with MPCs/MICs based cards. The "MPCE Type 3 3D" cards log the following messages and the forwarding ASIC on these cards may become stuck and start dropping all the packets that should have been forwarded. Aug 13 17:07:08.844 2013 rstMX480rmts2-lab-newton fpc2 XMCHIP(0): LI0: Received a parcel with more than 512B accompanying data Aug 13 17:07:08.845 2013 rstMX480rmts2-lab-newton fpc2 XMCHIP(0): LI1: Received a parcel with more than 512B accompanying data Aug 13 17:07:08.928 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) Congestion Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:09.009 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(4) Congestion Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:09.125 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(8) Congestion Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:09.240 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(12) Congestion Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:12.679 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) PPE\_0 Errors thread timeout error Aug 13 17:07:12.769 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) PPE\_1 Errors thread timeout error Aug 13 17:07:12.838 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) PPE\_2 Errors thread timeout error Aug 13 17:07:12.917 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) PPE\_3 Errors thread timeout error Aug 13 17:07:12.996 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) PPE\_4 Errors thread timeout error Aug 13 17:07:13.075 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0) PPE\_5 Errors thread timeout error Aug 13 17:07:30.728 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP:LUCHIP(8) Wedge Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:47.539 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP:LUCHIP(0) Wedge Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:47.662 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP:LUCHIP(4) Wedge Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:07:47.785 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP:LUCHIP(12) Wedge Detected, Active Zones f:f:f:f:f:f:f:f:f:f:f:f Aug 13 17:08:11.956 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.034 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(8): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.201 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.279 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(4): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.360 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(12): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.437 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(8): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.515 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.593 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(4): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.671 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(12): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.750 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(0): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.829 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(4): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.906 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(8): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:12.985 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(12): Secondary PPE 0 zone 1 timeout. Aug 13 17:08:13.064 2013 rstMX480rmts2-lab-newton fpc2 LUCHIP(4): Secondary PPE 0 zone 1 timeout. [PR876824: This issue has been resolved.](#)



- For a configuration with bridge domains containing aggregate interfaces, traffic whose destination address is broadcast, multicast, or unknown will not be load-balanced across the member links of such interfaces. Instead, all such traffic will be sent out a single link of the aggregate interface. With this PR change, load-balancing will always be applied to such configurations for traffic whose destination address is broadcast, multicast, or unknown. This change restores the functionality of older releases.

[PR888232: This issue has been resolved.](#)

- In VPLS environment, while deactivating/activating VPLS routing-instances, in rare conditions, routing protocol process (rpd) tries to free an already used route, then rpd process crashes with core files generated. [PR908856: This issue has been resolved.](#)

### **Layer 2 Ethernet Services**

- New knob is provided to set the prefix to compare requested ip and server address. Knob is configured as - [edit system services dhcp-local-server] #set requested-ip-network-match <0-31> For V6 [edit system services dhcp-local-server] #set dhcpv6 requested-ip-network-match <0-127> Default will be 8 for v4 and 16 for v6 (first terms). [PR872145: This issue has been resolved.](#)
- DHCPv6 Local Server implementation deletes the client on a reconfigure, so that client can reconfigure. DHCPv6 relay is not forwarding the Reply to the client and simply tearing the client down (generating a release to the server). [PR879904: This issue has been resolved.](#)
- When executing "show dhcp relay binding" command with high scales of bound subscribers and with several hundred renewing at a given time, DHCP drops the renew packets. [PR882834: This issue has been resolved.](#)

### **MPLS**

- The routing protocol process (rpd) might leak memory when there are MPLS LSP changes. The memory leak could eventually cause rpd process to crash. [PR847354: This issue has been resolved.](#)
- When BGP labeled-unicast route has BGP label as null and its indirect next-hop requires adding 2 or more labels, traffic using the BGP label may not be forwarded properly. [PR881571: This issue has been resolved.](#)
- To trigger the issue, there was a sequence of scheduling route-change and route-delete operations for the same LDP route. If the scheduling of a route-delete operation happens before the previously scheduled route-change operation is serviced, the crash will happen. The external event could be the Routing Engine switchover or link down. [PR912574: This issue has been resolved.](#)

### **Network Management and Monitoring**

- A memory leak in the cosd process is seen when both of the following conditions are met: - multiple OIDs from jnxCos MIB, that are under the same logical interface hierarchy, are queried in a single SNMP query sent to the device (i.e. in a single PDU) - either "per-unit-scheduler" or "hierarchical-scheduler" configured on the physical interface. The following messages will be logged when the cosd process exceeds 85% of its maximum usable memory: router-re0 /kernel: %KERN-5: Process (1457,cosd)

has exceeded 85% of RLIMIT\_DATA: used 1894060 KB Max 2097152 KB [PR893464](#): This issue has been resolved.

- In an IS-IS scenario, with traceoptions enabled under protocol IS-IS and syslog level set to debug under routing-options options for a router, if the router has two IS-IS neighbors which have the same router-id configured, after configuring the same ISO system-id on these two IS-IS neighbors, routing protocol process (rpd) on the router will crash with core files generated. [PR912812](#): This issue has been resolved.

### ***Platform and Infrastructure***

- RMOPD crash is due to sort of buffer overflow crash and library function being used improperly. It is not caused by RPM scaling. This issue happens randomly and hard to point out the specific trigger. [PR277900](#): This issue has been resolved.
- The "request system zeroize" command deletes the /var/db/scripts directory and all subdirectories but does not re-create them. The directories and subdirectories need to be manually re-created via the root shell and the correct permission set. [PR736478](#): This issue has been resolved.
- Junos OS 10.4R8 or later on MX Series platforms, L3VPN application using l3vpn-composite-nexthop when the indirect-next-hop configuration statement is added or removed, it might cause traffic drops affecting L3VPN flows. To recover from this condition all the l3vpn prefixes need to get removed and installed new into the forwarding-table, like clearing the bgp peers where the routes are learned from. [PR741646](#): This issue has been resolved.
- On TX Series system platforms, under very special corner case condition, non-enhanced FPCs might drop most of the traffic send into the fabric. TXP platforms are not exposed to this symptom. You will see lots of fabric drops reported via the "show pfe statistics traffic" or "show class-of-service fabric statistics" command. FPC affected needs to be restarted to recover from this condition. Sometimes you might also see the following error log reported in the syslog. LOG: Err] NFAB(1/1): RODR offset overflow count incremented (65) LOG: Err] CMALARM: Error (code: 542, type:Minor) encountered, cmalarm\_passive\_alarm\_signal LOG: Err] NFAB(1/1): PKTR ICELL signature error counter incremented [PR805682](#): This issue has been resolved.
- Under heavy traffic flow condition and non graceful FPC rebooting (i.e. temporary power failure on egress FPC) or SIBs getting automatic restarted to recovery from fabric connectivity issues, fabric ASIC on ingress T Series Enhanced Scalability (ES) FPC, can run into temporarily problematic status. This will cause temporary large delay on fabric traffic from T Series Enhanced Scalability (ES) FPCs to the egress FPC causing RODR offset overflow conditions and can have operational impact on transit traffic. This is only applicable to single chassis systems. The following syslog entries might get reported for many minutes. Sep 18 15:26:43 router fpc1 NFAB(1/0): RODR offset overflow count incremented (1) Sep 18 15:26:44 router fpc1 NFAB(1/0): RODR offset overflow count incremented (1) Sep 18 15:27:24 router fpc1 NFAB(1/0): RODR offset overflow count incremented (1) Sep 18 15:27:25 router fpc1 NFAB(1/1): RODR offset overflow count incremented (1) Sep 12 21:28:07.200 router fpc0 NFAB(0/1): %PFE-3: PKTR ICELL signature error counter incremented Sep 12 21:28:14.057 router fpc0 NFAB(0/1): %PFE-3: PKTR ICELL signature error counter incremented Sep 12 21:28:14.988 router fpc0 NFAB(0/1): %PFE-3: PKTR ICELL signature error counter

incremented Sep 12 21:28:15.989 router fpc0 NFAB(0/1): %PFE-3: PKTR ICELL signature error counter incremented Sep 12 21:29:19.989 router fpc0 NFAB(0/1): %PFE-3: PKTR ICELL signature error counter incremented [PR831743: This issue has been resolved.](#)

- Packets dropped with reject route are currently subjected to loopback filter processing on MPCs. As a result the packet dropped by a reject route may be seen in the output of "show firewall log". This behavior will be changed so that this traffic is no longer subjected to loopback filter processing to bring it in line with other line cards. [PR858511: This issue has been resolved.](#)
- With an interface-specific filter contains a percentage policer configured on several interfaces, when the shaping rate of an interface changed, the percentage policer instances of the filter applied on that interface need to be updated. If FPC restarts when policer instances are being updated, an interface-specific filter instance might not be instantiated in hardware, causing FPC to dereference a NULL pointer, then FPC crashes with core files generated. [PR874923: This issue has been resolved.](#)
- There are two symptoms covered this issue: If there is a mix of high and low priority fabric traffic as can be seen by checking "show class-of-service fabric statistics", the following error messages can be seen when there are bursts of high priority fabric traffic, while low priority fabric traffic is present : May 6 14:58:41 routename-re0 fpc1 MQCHIP(0) FI Reorder cell timeout May 6 14:58:41 routename-re0 fpc1 MQCHIP(0) FI Cell underflow at the state stage A second symptom with this mix of low and high priority fabric traffic present; if an FPC that is the recipient of this high and low priority fabric traffic restarts, it is possible for the ingress FPC forwarding ASIC to lock up. In this case the following log message might be simultaneously logged :- Jun 5 13:46:50 router fpc4 MQCHIP(0) CPQ Queue underrun error, Qsys1 Queue 42 Jun 5 13:46:50 router fpc4 MQCHIP(0) CPQ Freecnt nearing empty error, Qsys mask 0x2 [PR877123: This issue has been resolved.](#)
- This is a regression issue introduced by the fix of PR801982, which causes DOM MIB values for SFP+ "rx power" related statistics are incorrect. Please note that XFP is not affected. [PR878843: This issue has been resolved.](#)
- If interface flaps of a bridge-domain with igmp-snooping enabled or multicast snooping routes are pruned due to Designated Router changes, LUCHIP might report traps and EDMEM read errors. These conditions are transient and only seen once the system is operating with enhanced-ip mode. [PR879158: This issue has been resolved.](#)
- For MX Series based FPC, on PHP->PE link, custom MPLS MTU allows more than configured size. [PR879427: This issue has been resolved.](#)
- Deactive/delete AE interface when route is flapping might cause the MX Series based Packet Forwarding Engine to crash. [PR884837: This issue has been resolved.](#)
- In l2circuit connection scenario, when the lchip based FPC interconnect with MX Series based FPC, PPP-CCC l2circuit connection will drop the small packets with Ethernet length error. [PR887098: This issue has been resolved.](#)
- In L2VPN scenario, on the PE router, if the encapsulation of the PE-CE interface is vlan-ccc and there is a COS filter under the interface, when the interface flaps, it can cause all the traffic to different sites via different outgoing interfaces to be forwarded incorrectly through one of the interfaces. Meantime, when manually flap the

label-switched paths (LSPs) on the router after the problem occurred, the traffic is forwarded incorrectly still but only the egress interface will change to other one. The way to resolve the problem is manually clearing the LSPs on the PE router. [PR887838: This issue has been resolved.](#)

- It is observed that in the setup route nexthop for destination of collector's IP address was of type indexed nexthop. [PR889884: This issue has been resolved.](#)
- Because of the hardware limit, the feature "maximum-labels" on FPC can't exceed 3. Whenever maximum mpls label is configured as 4 or 5 on unsupported FPC, the LDP/RSVP session will go down and cause MPLS traffic black hole for couple of minutes. This dark window will remain till the unicast next hops are installed and attached to the egress interface where the label has been configured. After that MPLS traffic will resume. [PR890992: This issue has been resolved.](#)
- Traffic may be affected after performing an offline/online sequence on the PIC in a T4000 system. This issue is usually seen when the event is performed on PICs carried in a Type 5 FPC. [PR892548: This issue has been resolved.](#)
- When a filter/fw config is modified, poisoned next-hops (log message Packet Forwarding Engine: Detected error nexthop) are reported and an automated jsim is performed on the affected packets. This is happening on the Packet Forwarding Engines with two jtree segments and the issue is transitory. [PR897107: This issue has been resolved.](#)
- In MX-VC setup using virtual-switch instance type, there can be scenarios where the outer vlan-tag of PPPoE/PADI packets on egress can be stripped off when ingress interface is a LAG with two member links spread across the two chassis members. [PR905667: This issue has been resolved.](#)
- Command "show ddos-protection protocols" doesn't report correct Arrival and Max arrival pps rates. One bit of rate value at the Packet Forwarding Engine is incorrectly set which results in a wrong ddos rate value. [PR908803: This issue has been resolved.](#)

### ***Routing Policy and Firewall Filters***

- Install-nexthop lsp-regex does not work as expected when multiple recursive routes share same protocol next hop having different export policy with regular expression option. Route is not updated with correct export forwarding nexthop as same nexthop select handle is calculated for any set of configured export policy with "install-nexthop lspregx" option. [PR863341: This issue has been resolved.](#)

### ***Routing Protocols***

- In some scenarios MVPN-routes with same RD:Prefix may get generated from multiple-VRFs on a PE-router. When such a PE-router is not a MVPN-RR and has no MVPN-EBGP peers, it is possible that the core-network may lose the MVPN-route because of an erroneous MVPN-withdrawal sent by the PE because of the MVPN-route getting deleted from one of the PE VRFs; even if there are other-VRFs on the PE still advertising the route. [PR698493: This issue has been resolved.](#)
- BFD triggered local-repair(RLI9007) not initiating immediately. RLI 9007 is applicable from 12.2 onwards. [PR825283: This issue has been resolved.](#)

- In BGP scenario, the initial peer flaps and goes down, then a new peer is established which might cause an rpd core file. [PR840652: This issue has been resolved.](#)
- When configuring CAC for a physical interface, the software might enable CAC for unit 0 on that interface, but might not be able to delete it when the configuration is removed. [PR850578: This issue has been resolved.](#)
- If nonstop active routing (NSR) is enabled and fxp0.0 is configured under IS-IS as "disable" and not "point-to-point", after the Routing Engine switchover the backup Routing Engine generates a pseudonode for fxp0.0 treating it as a regular LAN interface. This causes the backup Routing Engine to generate an IS-neighbor to itself. [PR861743: This issue has been resolved.](#)
- In an IS-IS scenario, when a large number of routes are distributed into IS-IS, IS-IS overload bit will be set due to maximum LSP fragment exhaustion. This is correct. Then delete the IS-IS export policy, after that, the IS-IS overload bit should be cleared. But the number of exported prefix might be incorrect even though the number of export prefix is zero actually. This can cause overload bit to be set always. This is because local-data for prefixes is not freed up and leads to some memory leak. [PR874015: This issue has been resolved.](#)
- The remote discriminator is not reinitialized after bfd session state moves to down (with diagnostic code: control detection time expired) as per RFC 5880 requirement. [PR889970: This issue has been resolved.](#)
- The downstream PE router's RPF\_neighbor(S) on the MDT reverts back to mRIB.next\_hop(S) rather than the Assert(S,G)Winner when their PPT expires. Bug identified in the code and is fixed. [PR896898: This issue has been resolved.](#)
- In PIM dense mode, if the Assert loser router receives a join/prune (S,G) message with upstream neighbor is the loser router, it should send a Assert(S,G) on the receiving interface to initiate a new Assert negotiation to correct the downstream router's RPF neighbor, but our device will not. This PR has solved the issue. [PR898158: This issue has been resolved.](#)

### **Services Applications**

- Memory leak in key management daemon (kmd) causes some IPsec VPN tunnels to be dropped and don't get re-negotiated for over 10 minutes. Before issue happens, the following logs could be observed: /kernel: Process (1466,kmd) attempted to exceed RLIMIT\_DATA: attempted 131080 KB Max 131072 KB /kernel: Process (1466,kmd) has exceeded 85% of RLIMIT\_DATA: used 132008 KB Max 131072 KB [PR814156: This issue has been resolved.](#)
- This issue is seen when two l2tp users get connected to same routing-instance and they get same framed routes. When last connected user disconnects this issue can be seen. [PR832034: This issue has been resolved.](#)
- Enabling KMD traceoptions (with level set to warning, all/verbose/notice/info) results in KMD core during rekey procedure. [PR856499: This issue has been resolved.](#)
- MIB module in file "mib-jnx-sp.txt" contains a coding error, which may lead to a loop. [PR866166: This issue has been resolved.](#)

- If RSP1 and RSP10 interfaces are configured on the same box issuing the "request interface switchover rs1" or "request interface revert rsp1" causes both RSP1 and RSP10 to switchover or revert. [PR877569: This issue has been resolved.](#)
- In a CGNAT environment when sp interfaces, which are underlying rsp interface, are present in the configuration, sp interfaces service-options may wrongly overwrite rsp interfaces service-options and syslog stopped working and inactivity-timeout values were reset to the default values. [PR881792: This issue has been resolved.](#)
- The jl2tpd process generates a core file as follows:  
"././src/bsd/lib/libc/stdlib/abort.c:69." [PR887662: This issue has been resolved.](#)
- The jpppd crash on LNS happened because the size of the udp based l2tp packet exceeded the buffer length available. The modification was done to discard the packet instead of creating core. [PR888691: This issue has been resolved.](#)
- SIP ALG - Service PIC might crash when SIP flows are cleared. [PR890193: This issue has been resolved.](#)
- Output interface' shown as 'Unknown' under show services accounting flow-detail.issue has been analysed RCA;-At the time when a flow is created in PIC memory, if the route to the destination IP(in the flow) is not known, we set a flag indicating that there is no route to Destination IP in the flow structure. When the flows are queried using "show service accounting flow-detail" picinfo daemon inspects this flag for each flow and prints the Output interface as "Unknown" if this flag is set. Now, after route record for that flow is downloaded to the Service PIC, the flow structure is updated to reflect the corresponding output interface, but, the above flag is NOT UNSET. So, picinfo daemon continues to print the output interface as "unknown" whenever "show services accounting flow-detail" is executed. [PR890324: This issue has been resolved.](#)
- L2TP session on MS-PIC may fail and following error is observed  
"L2TPD\_RADIUS\_SERVER\_NOT\_FOUND" after a test access profile <ppp-profile> is issued. [PR898872: This issue has been resolved.](#)
- When the 'learn-sip-register' knob is enabled for the SIP ALG (it is by default), for a SIP request in slow path implicitly denied by the firewall or NAT rules, a look up is done to see if the SIP request has a target that corresponds to any current registration state, in which case the corresponding reverse flows get created. While service PIC creating the corresponding reverse flows, an internal error may occur, causing service PIC to crash and create a core file. [PR899195: This issue has been resolved.](#)
- In Carrier Grade Network Address Translation (CGNAT) environment, if memory utilization of MS-DPC/service PIC are in the yellow zone and they are configured with cgn-pic knob, there can be a race condition where there are two timers created for the same flow and during the timer processing, the MS-DPC/service PIC may experience a crash and create a core file. [PR901795: This issue has been resolved.](#)
- In some cases rtsp data flows will be left without clean up when rtsp master flow closes. This will cause some conversation data flows left on router with very huge timeout value. [PR909091: This issue has been resolved.](#)

## VPNs

- VPLS traffic gets flooded back over the ingress interface on the local PE as the split-horizon gets disabled upon interface flap. [PR818926: This issue has been resolved.](#)
- In this release Ngen-MVPN does not support NSR. But the commit check when Ngen-MVPN and NSR are configured does not fail. In previous releases this commit would fail. The commit check not failing for this configuration is planned to be fixed in release 12.3 R4. In Release 12.3 R3 config with NSR and Ngen-MVPN configuration should not be committed. Doing this commit can lead to routing application crashes (like PR 864439) as it is an unsupported feature. [PR827519: This issue has been resolved.](#)
- In case graceful-restart (GR) is enabled for BGP-MVPN, after rpd restart on ingress PE, PIM never installs PIM (S,G)route IF the CE1->PE1 BGP connection comes up after the PE1->RR BGP connection. As a result, Multicast flows stop. Graceful-restart with BGP-MVPN is not a supported feature. [PR872009: This issue has been resolved.](#)
- If SNMP "get" tries to retrieve local and direct routes from mplsL3VpnVrfRteTable, they are not found. SNMP walk does walk the local and direct routes. [PR874365: This issue has been resolved.](#)
- In VPLS scenario, if CE facing interface is aggregated Ethernet with multiple member ports (more than two members), BUM (broadcast, unicast unknown, and multicast) traffic from MPLS core will be replicated on all child links of aggregated Ethernet interface and BUM from CE will be replicated at sending out from MPLS core facing interface. The problem is specific to M10i and M7i routers running with I-chip based CFEB. [PR880422: This issue has been resolved.](#)

## Release 12.2R5

### Class of Service (CoS)

- A few memory leaks have been fixed in the class of service daemon. [PR811613](#)
- This cosmetic issue is specific of 3D linecards, based on MX Series routers with MPCs/MICs chipset. In these cards, the logical interfaces with family mpls do not have any EXP rewrite rule applied by default. In other words, EXP value is copied from the previous codepoints: for example, from IP Precedence in IPv4->MPLS next hops. However, the command "show class-of-service interface" still shows the exp-default rule as if it was applied (in fact, it isn't): user@router> show class-of-service interface ge-2/3/1.204 | match rewrite Rewrite exp-default exp (mpls-any) 33 . [PR824791.](#)
- When 'scheduler-map-chassis derived' configuration is used under class-of-service, interface related configuration changes can lead to cosd process crash. [PR863734](#)

### Forwarding and Sampling

- Possibility of duplicate packets when sampling and interface-style nat are configured. [PR861984](#)
- If there is distributed Bidirectional Forwarding Detection (BFD) running on Aggregated interface and a firewall filter is configured on loopback interface (lo0), the lo0 will bind an implicit filter, after FPC restarts or Routing Engine switchover, the next hop of the implicit filter is not updated with the corresponding link word to point to CLI filter,



causing the CLI filter to be not executed. To resolve the issue, deactivate the firewall filter under loopback interface and then activate it again. [PR864665](#)

### **General Routing**

- The 'RL-dropped' lines of "> show interfaces queue" will be missing when the PIC is bounced. [PR749283](#)
- Changing static route with qualified-next-hop and order option to next-hop option results in static route missing from route table. We need to restart routing process to see the route again. [PR830634](#)
- It is possible that RPD's higher priority tasks (HPTs) are scheduled to run such that lower priority tasks (LPTs) may not be able to complete until HPTs are completed. This problem has been resolved and fix available in the following releases: 12.3R3 12.2R5 13.1R1 13.2R1 11.4R8 12.3R3 12.1R7 11.4R7-S2 [PR836197](#)
- After a Routing Engine switchover with graceful Routing Engine switchover (GRES) enabled, and then deactivate and activate a routing-instance, 4xOC48 IQE PIC might reboot unexpectedly. This is caused by a problem in channel allocation for the 4xOC48 PIC logical interfaces in kernel. [PR841822](#)
- When MX Series router running with DPC is upgraded by unified ISSU, some of interface may show incorrect input packet/byte count. And the incorrect count is also seen to the related interface MIB. The value will be a large number. Physical interface: xe-3/1/0, Enabled, Physical link is Up Interface index: 138, SNMP ifIndex: 5449, Generation: 141 Link-level type: Ethernet, MTU: 1514, LAN-PHY mode, Speed: 10Gbps, BPDU Error: None, Loopback: Local, Source filtering: Disabled, Flow control: Enabled Device flags : Present Running Loop-Detected Interface flags: SNMP-Traps Internal: 0x4000 Link flags : None CoS queues : 8 supported, 8 maximum usable queues Hold-times : Up 0 ms, Down 0 ms Current address: 00:24:dc:9c:7c:30, Hardware address: 00:24:dc:9c:7c:30 Last flapped : 2013-01-13 14:36:25 JST (02:07:52 ago) Statistics last cleared: Never Traffic statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps ~snip~ Logical interface xe-3/1/0.0 (Index 196614) (SNMP ifIndex 5450) (Generation 140) Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Traffic statistics: Input bytes : 3867797326912475 Output bytes : 0 Input packets: 15108583308733 Output packets: 0 Local statistics: Input bytes : 0 Output bytes : 0 Input packets: 0 Output packets: 0 Transit statistics: Input bytes : 3867797326912475 0 bps Output bytes : 0 0 bps Input packets: 15108583308733 0 pps Output packets: 0 0 pps Protocol inet, MTU: 1500, Generation: 160, Route table: 0 Flags: Sendbroadcast-pkt-to-re Addresses, Flags: Is-Preferred Is-Primary Destination: 10.3.1/24, Local: 10.3.1.1, Broadcast: 10.3.1.255, Generation: 141 Protocol multiservice, MTU: Unlimited, Generation: 161, Route table: 0 Policer: Input: \_\_default\_arp\_policer\_\_ gladiolus:Desktop\$ grep .5449 mib\_value\_after\_issu.txt ifName.5449 = xe-3/1/0 ifInMulticastPkts.5449 = 0 ifInBroadcastPkts.5449 = 0 ifOutMulticastPkts.5449 = 0 ifOutBroadcastPkts.5449 = 0 ifHCInOctets.5449 = 3867797326912475 ifHCInUcastPkts.5449 = 0 ifHCInMulticastPkts.5449 = 0 ifHCInBroadcastPkts.5449 = 0 ifHCOctets.5449 = 0 ifHCOUcastPkts.5449 = 0 ifHCOUmulticastPkts.5449 = 0 ifHCOUbroadcastPkts.5449 = 0 gladiolus:Desktop\$ grep .5450 mib\_value\_after\_issu.txt ifName.5450 = xe-3/1/0.0 ifInMulticastPkts.5450 = 0 ifInBroadcastPkts.5450 = 0



ifOutMulticastPkts.5450 = 0 ifOutBroadcastPkts.5450 = 0 ifHCInOctets.5450 = 3867797326912475 ifHCInUcastPkts.5450 = 15108583308733  
 ifHCInMulticastPkts.5450 = 0 ifHCInBroadcastPkts.5450 = 0 ifHCOctets.5450 = 0 ifHCOUcastPkts.5450 = 0 ifHCOMulticastPkts.5450 = 0  
 ifHCOBroadcastPkts.5450 = 0. [PR847106](#)

- It is possible for RPD core when the following conditions are met: - VRF with multipath knob configured - static routes with next-hops which are indirect type and needs further resolution - the numerically lowest (smallest IP) next-hop of indirect type becomes unreachable RPD core is NOT triggered in either of the following scenarios: - no multipath under VRF - if there is no static route entry - static route whose next-hops are indirect type requiring further resolution multipath under VRF is supported only for BGP configurations. multipath in other conditions are not supported, and a bug in this detection phase is fixed in this PR. [PR847214](#)
- In certain graceful Routing Engine switchover (GRES) scenarios, with IPv6 address configured on at least two interfaces, Solicited node multicast addresses (SNMA) and link local addresses with same prefix might be created on the two interfaces. There is a possibility that there could be inconsistency in the Next Hop database between Master and Backup Routing Engines. When the Backup becomes Master in these scenarios, it'll try to program the Packet Forwarding Engines with the bad Next Hop data. This may cause undesired forwarding behavior on the Packet Forwarding Engines. [PR850625](#)
- Routing Engine may kernel panic crash after software upgrade to Junos OS including the fix of PR831233. [PR851086](#)
- Ptsf failed to append policy with multi-rules since 'msg over size limit'. [PR852224](#)
- FPC or PIC connects to Routing Engine Kernel for the first time when it comes up or reconnects during connection trip. After the connection is established with Routing Engine, if FPC/PIC does not respond kernel for 300 seconds, a timer is triggered to disconnect Routing Engine from FPC/PIC. In a particular race condition between kernel processing received data on the connection and the fired timer trying to close the connection, kernel crashes and generates core file. FPC/PIC's slow response may be attributed to high traffic or a faulty hardware. Before kernel crash, the following logs could be seen: fpc3 LCHIP(3): 1 new Lin SIF ins eoep errors fpc3 LIN(3): PIC HSR is not OK, LCHIP(3) <- PIC 3 HSR 1. [PR853296](#)
- ATM MIC back-to-back, to many logical interfaces(more than 8k) might cause certain logical interfaces down. [PR859165](#)
- In a virtual chassis environment in the event power is loss on the Master virtual chassis the standby chassis has potential to experience slot resets during transition period. [PR859717](#)
- Multiple clksyncd core-dump may be seen after unified ISSU upgrade on the chassis. [PR861676](#)
- When a prefix next-hop address resolution requires a recursive lookup, the next-hop might not be updated correctly after an egress interface is disabled. [PR862989](#)
- When using BGP Flow Spec with rate-limit option, even though the value is in Bytes/second, the value being programmed is in bits/second. [PR864496](#)

- Output of "show subscribers physical-interface aex" displays multiple AE links. [PR864555](#)
- Configuration of Container Interfaces for APS on MX Series FPCs is not allowed since Junos OS 12.1. If this feature is needed on MX Series legacy FPCs, use a release with this PR fixed. [PR869192](#)
- When configuration stanza: [protocols router-advertisement] starts as: ## ## inactive: protocols router-advertisement ## interface ge-0/0/1.1 { virtual-router-only; } Then perform the following actions: Step 1 - activate protocols router-advertisement Step 2 - deactivate protocols router-advertisement interface ge-0/0/1.1 Step 3 - set protocols router-advertisement interface ge-0/0/1.2 After issuing "commit check", there are no problems. But after issuing "commit", routing protocol process (rpd) crashes and generates a core file with following logs: rpd[1422]: RPD\_RA\_CFG\_UNKNOWN\_ACTION: Unknown configuration action 3 received. [PR871359](#)
- Under high scale, expiry of a Kernel side reconnect timer would cause it to send a non servicable message to the Packet Forwarding Engine (asking the line cards to restart and re-sync since reconnect failed) Since there is no ack- to this Kernel message, Kernel thought it sent the message and untoggles the GRES flag. The Packet Forwarding Engine wasn't expecting anything so it continued along. The EFFECT: The system is permanently not ready for GRES... CLI GRES check will always report: [cmd] request chassis routing-engine master switch check Apr 14 19:03:13 [INFO ] warning: Standby Routing Engine is not ready for graceful switchover. [PR873679](#)
- The default setting for the sysctl "net.pfe.relayg\_merge\_enabled" is 0 (off). This results in a support limit of 16 line cards within the VC. Even with the group merge disabled, line-cards may have been grouped at system start-up only presenting an issue after they restart. [PR874791](#)

### **Infrastructure**

- The root cause of the problem was IFADDR change in VRRP context was not replicated to GRESS backup. [PR790485](#)
- Kernel fails to generate ICMP ttl expired when IP packet len is a multiple of 256. [PR829567](#)
- Delay in bringing ONLINE an FPC after it is inserted into the chassis. [PR853304](#)
- In a scenario with scaling routes existing (e.g. 54k BGP routes), while these routes flapping, in a rare case, the TCP connection between Routing Engine and FPC is mistakenly enabling re-transmit timer for pure ACK's which is causing the FPC to reboot. [PR858489](#)
- With nonstop active routing (NSR) enabled, while performing graceful Routing Engine switchover (GRES), Junos OS fails to restore BGP peers' TCP connections on the new master Routing Engine's replicated socket due to it is not able to find the BGP peer address's route, causing BGP peers to flap with following logs: /kernel: jsr\_sdrL\_merge: PSRM merge failed 65 rpd[xx]: RPD\_BGP\_NEIGHBOR\_STATE\_CHANGED: BGP peer a.b.c.d (Internal AS X) changed state from Established to Idle (event TcpSocketReplicationError). [PR862796](#)

- When a sonet interface with PPP encapsulation is used as forwarding next hop for the IPv6 remote router loopback address on IPv6 BGP sessions, if the sonet link is down, the IPv6 BGP session might flap at same time although there is valid route via other interface. [PR863462](#)
- After enabling firewall filter of IPv6 on Aggregated Ethernet (AE) interface to block Micro BFD Packets (Dst Port 6784), kernel crashes continually on Master and Backup Routing Engine due to double free of memory. [PR864112](#)
- IPv6 Neighbor discovery(ND) failed after multiple GRES. Nexthop getting stuck in hold state forever. We also see that the neighbor state is in NO\_STATE and it is on ND timer queue. In this condition, on ND timer expiry it never sends neighbor solicitation (NS) out and it never transitions to known ND states. Use "show route forwarding-table" CLI command to see the result of IPv6 route in hold state. root@ABC> show route forwarding-table Destination Type RtRef Next hop Type Index NhRef Netif 1234::56 /128 dest 0 1234::56 hold 1902 1 irb.5678 Use "show ipv6 neighbors" CLI command to see the result of IPv6 ND state in NO\_STATE. root@ABC> show ipv6 neighbors IPv6 Address Linklayer Address State Exp Rtr Secure Interface 1234::56 none nostate 0 no no irb.5678. [PR864133](#)

### ***Interfaces and Chassis***

- There can be a mismatch between the ifIndex value on IF-MIB-ifName and the ifIndex value on SONET-APS-MIB-apsMapGroupName and apsMapEntry. [PR771877](#)
- This issue is specific to the M120 hardware since there are two independent FRU's from where the PIC needs to be detached/attached. This IPC messages go out-of-order due to the additional control-plane messages related to routing-change as a result of PIC restart which happens in this case due to the buffer configuration change. When PIC needs to be detached and at the same time there is still a lot of protocol information which should be processed as well, the ?detached? messages will NOT be able to be delivered in time. After PIC restarts it requests to be attached again but obviously this action failed because from other FRUs perspective the PIC has NOT been detached at all. [PR773081](#)
- Hash Key configuration not programmed in Packet Forwarding Engine correctly after system reboot. [PR818035](#)
- Faulty SCG causes continuous interrupts to HCFPC making its CPU Utilization 100% and unusable for any service. As a fix the monitoring mode for the SCG is changed to polling status of SCG device rather than interrupts based awake and monitoring system. [PR827489](#)
- Internal interfaces used by the Routing Engine to communicate with all of the Packet Forwarding Engines and the standby Routing Engine can sometimes fail auto-negotiation after a reboot. The interfaces can be recovered using the "ifconfig up" command in the shell, but there is now an automatic detection and recovery mechanism to restore the failed interfaces. [PR829521](#)
- Removing IP address on ATM interface after adding another IP address from the common subnet can lead to a race condition. New IP address configured on the interface still referring to shared broadcast-nexthop. Then when TCP/IP accesses this broadcast-nexthop, kernel panic may happen. [PR833015](#)

- When packet has to be forwarded over NH topology unilist->indirect-indexed and when the packet size is greater than egress interface MTU with DF set, then we may log the following message and not send the message back to source indicating "frag needed and DF set". fpc0 NH: Can not find logical interface for nh 1048590 fpc0 NH(nh\_get\_mtu\_iff) : get unilist mtu failed. [PR844987](#)
- Whenever tunnel interface -pe/-pd got created using the MS-DPC instead of the MPC, it will not be able to process register messages. Because of MPC and MS-DPC have different multicast architectures and they are incompatible if chassis is configured in "enhanced-ip" mode, this issue will be seen. Necessary changes have been made to code so that these interfaces will not be created on MS-DPC. [PR853995](#)
- SDG : After rebooting both Routing Engines together, the FPCs and MS-DPCs may come online, go offline (with "Chassis connection dropped" and "Chassis Manager terminated" error messages) and come back online again automatically. This issue is seen only when both Routing Engines are rebooting at once. There is exactly one additional reboot of the FPCs when this happens, and the FPCs come back up online, and system stabilized by itself within 2 to 3 additional minutes [PR/854519: This issue has been resolved in 12.1X43.3]. [PR854519](#)
- Routing Engine reset with Fatal trap 12: page fault while in kernel mode. [PR855317](#)
- Multilink Frame-relay (MLFR) stuck in ready state after restarting FPC and then graceful Routing Engine switchover (some of the MLFR bundles will show "ready" although the interfaces are in up/up state which causes data loss). [PR857648](#)
- The backup Routing Engine might log the following often in chassisd: Feb 17 12:40:01 CB:1 need not to sync information Feb 17 12:40:21 CB:1 need not to sync information Feb 17 12:40:41 CB:1 need not to sync information Feb 17 12:41:01 CB:1 need not to sync information This is a harmless message that can be ignored. [PR857698](#)
- Interface hold-time-down is not working properly for PIC type 10x10GE(LAN/WAN) SFPP. [PR859102](#)
- When a pppoe subscriber sends a 'LCP Configure-Request' message with configuration option 'Authentication Protocol PAP', MX Series BNG responds with 'LCP Configure-Ack', instead of rejecting it with 'LCP Configure-Reject'. After sending LCP 'Configure-Ack', BNG continues by sending 'PAP Authenticate-Request', with blank 'Peer-id' and 'Password'. This makes MX Series BNG behave like a client on PPP Session. Since MX Series BNG is always supposed to have a Server role in PPP Session, it must respond with LCP Configure-Reject, whenever it receives LCP Configure-Requests with 'Authentication Protocol' option. [PR860089](#)
- Unified ISSU does not support VRRP. [PR862052](#)
- snmpwalk of "jnxPPPoEIfLockoutTable" did not capture pppoe locked out clients. [PR869024](#)

- Injecting Enhanced RDI-P(G1 bit5-7:0x2 Payload defect) alarm to a MPC 10GbE WAN-PHY interface causes RDI\_P and LCD-PAIS-V alarm on messages. This is due to string typo. RDI\_P and LCD-P should be printed on messages. [PR872133](#)
- Both VRRP routers keep backup-backup state until "startup-silent-period" expires if both "startup-silent-period" and "delegate-processing" are configured. The fix is available in 12.1R7, 12.2R5, 12.3R3, 13.1R3, 13.2R1 and later release. [PR873488](#)

### Layer 2 Features

- When VPLS is configured with GRES, the backup Routing Engine responds to certain route replication requests by simulating address learning. If the route being replicated is associated with an LSI or VT interface, the address learning code references a special LSI or VT nexthop. Thus, there is a dependency between that route and that nexthop. This fix is to explicitly enforce this ifstate dependency, ensuring that the special nexthop is seen by the peer before the route. [PR867929](#)

### Layer 2 Ethernet Services

- jdhcpd interface traceoptions are not saved to the default log file jdhcpd and require an explicit file name. [PR823129](#)
- DHCPv6 fails for clients using DUID type 2 (Vendor-assigned unique ID), the software was using the DUID to extract MAC address information. This behavior is fixed and tested. [PR838404](#)
- MXVC-DHCP bindings stuck in a "RELEASE(RELAY\_STATE\_WAIT\_AUTH\_REQ\_RELEASE" state. [PR850187](#)
- In DHCP subscriber management environment, while DHCP subscribers login, in rare conditions, system calls of these subscribers fail, due to only on success does system free the memory, resulting in a memory leak for the jdhcpd process. If memory usage of jdhcpd process goes to its limit, no new DHCP subscribers can login. When issue happens, high weighted CPU usage of jdhcpd process and following logs could be observed. /kernel: %KERN-5: Process (31403,jdhcpd) has exceeded 85% of RLIMIT\_DATA: used 2825132 KB Max 3145728 KB jdhcpd: %USER-3-DH\_SVC\_RTsock\_FAILURE: Error with rtsock: rtslib: ERROR Failed to allocate new block of size 16384 jdhcpd: %USER-3-DH\_SVC\_RTsock\_FAILURE: Error with rtsock: rtslib: ERROR Failed to allocate new block of size 16384 jdhcpd: %USER-3-DH\_SVC\_RTsock\_FAILURE: Error with rtsock: rtslib: ERROR Allocation Failure for (16384) bytes authd[1822]: %DAEMON-3: .././../src/junos/usr.sbin/authd/plugin/radius/authd\_plugin\_radius\_module.cc:1090 Failed to get SDB snapshot for session-id:3549005. [PR856024](#)
- In DHCP subscriber management environment, with scaled DHCP subscribers login, after executing "clear dhcp relay binding all/interface" or "clear dhcp server binding all/interface", new subscribers login are delayed, and it shows high CPU usage for a while. [PR857006](#)
- DHCP relay functionality over IRB performing dhcp v4 relay functionality, and configured with both inet and inet6 address families. The removal of the IPv6 address family configuration from the IRB can cause the IPv4 dhcp relay functionality on that IRB to break. This happens regardless of whether the 'family inet6' is configured directly under

the IRB or applied through a 'apply-group' configuration. In versions that do not have the fix for this PR, the workarounds to get the dhcp relay functionality working again over the IRB are *either* of the following 1) deactivate/activate the IRB configuration 2) Restart dhcp daemon using the following command `user@host> restart dhcp-service`  
Note: This workaround has to be applied everytime any configuration change (as explained in the trigger) is applied that could potentially get the dhcp-relay functionality to break. [PR870543](#)

- When IPv6 is configured on integrated routing and bridging (IRB) interfaces that have AE interfaces as child links, after GRES was enabled and after one child link failure or removal, the kernel crashed. [PR878470](#)

### ***MPLS***

- The cleanup procedures may leave transient inconsistent references when the interface address of an MPLS enabled GRE or IPIP tunnel is being deleted or the action taken implies an internal reconfiguration of the interface address (for example MTU change). During these periods, if these references are being reused by a particular task, the kernel may report an invalid memory access and restart. [PR844790](#)
- The routing protocol process (rpd) might leak memory when there are MPLS LSP changes. The memory leak could eventually cause rpd process to crash. [PR847354](#)
- Unsupported feature warning missing for mLDP+NSR while doing unified ISSU. [PR849178](#)
- RPD core observed on backup Routing Engine with `rsdp_mirror_telink_attempt_resolve`. [PR859602](#)
- ASBR might not rewrite EXP correctly for egress MPLS packets on the Inter-AS link for the eBGP-LU LSP if the eBGP session is a multihop BGP session. [PR864914](#)
- Apply group with session parameters will not work for LDP protocol from 12.2 release onwards without the fix for this PR. This is due to re-organization of 'ldp session' configuration during 12.2 development. [PR868945](#)

### ***Network Management and Monitoring***

- Flapping interfaces in combination with restarting `chassisd`/`dcd`/`mib2d` daemons and other abnormal scenarios coupled with SNMP polling create race conditions in `mib2d`. This results in a `mib2d` core. [PR812019](#)
- Under certain conditions, duplicate SNMP indexes might be assigned to different interfaces by kernel to `mib2d` (Management Information Base II daemon). This might cause `mib2d` and other daemons such as `lacpd` (LACP daemon) to crash and generate core files. [PR836823](#)

### ***Platform and Infrastructure***

- XML tags for `get-software-information` output missing some elements of new Junos OS service release naming convention. [PR783653](#)
- Tunnel services (MT-x/x) using MPC on PE installs (S,G) route on receiving IGMP report. [PR821893](#)

- In a race condition where multiple interrupts are asserted, timer tick may not get well handled and remain asserted. This caused FEB panic and core. [PR828496](#)
- When you archive a file using the file-archive rpc option, the following error is displayed: Operation allowed only from CLI. [PR831865](#)
- Due to a bug in logical interface localization, a DPC restart/offline may cause a removal of legitimate CCC routes on other DPC's. This can also be triggered by removal of an unrelated family CCC logical unit. [PR835216](#)
- FPC core dump with the feature copy-plp-all enabled when add link to existing AE interface, which is part of downstream interface list of a multicast route. [PR842046](#)
- When a junoscript <get-configuration> RPC query, by default the query is done on candidate DB, a MGD process is spawned to handle this request. Now at the same time via another session if the configuration is deleted it is possible for the above spawned MGD process performing the junoscript query to crash. MGD process crashes while accessing a NULL parent which contained an object previously which was deleted. The fix addresses this by not exporting the object which has no parent. [PR844795](#)
- mlfr/mlppp interfaces are not reachable after restart FPC (primary MSPIC) followed by deactivate and activate R.I or GRES followed by deactivate and activate R.I. This is because link FPC does not have the interfaces programmed towards the bundle. [PR847278](#)
- If routing-instance is popping the mpls label through vt tunnel interface and the egress interface MTU of the vrf needs fragmentation and the dont-fragment bit is set in the ipv4 header, the egress vrf interface might stop forwarding traffic. The following syslog message will be reported fpc4 LCHIP(3): 1 new errors in LSIF To recover from this condition you can either bring the interface down via disable knob or deactivate/activate the interface from the configuration. The following platforms are exposed to this condition: M320 (excluding E3 FPCs), T/TX systems (excluding ES FPCs and FPC Type 5). [PR854806](#)
- In the T4000 Type 5 FPC platform, aperture management can lead to a collision between the sched tick timer and asic driver interrupt handlers, which will result in FPC crashes. [PR857167](#)
- On MX Series routers, with some logical interfaces of an aggregated Ethernet (AE) interface attached to a bridge-domain and LACP is enabled on the AE interface, after disabling/enabling or removing/adding one or more member links of the AE interface, because the receive channel of the AE interface is closed when LACP state is down, traffic loss might be observed for several seconds. [PR858124](#)
- In IPFIX context: 1. In an IPv6 single stack environment, when exporting Data and Template records for family IPv6, the Template records sequence number is not initialized and is always == 0 for all records. This is because the Template sequence numbers are blindly copied from family IPv4 and if this is not configured for IPFIX, then the Sequence Number is always 0. 2. In an IPv4 + IPv6 dual stack environment, since the Template records sequence numbers will be identical for both families, we will get Data and Template records sequence numbers being interleaved when exported. This could confuse the Flow Collector and mislead it into reporting random missing flows. [PR859169](#)



- BOOTP request packets might get dropped because of the DDOS protection feature in MX Series routers with MICs/MPCs. In this case, the bootp packets are coming with 1 byte option. So the length of bootp become 241 which is larger than 240. Then the Packet Forwarding Engine will identify it not as BOOTP as per the current DDOS algorithm, and tries to parse it as DHCP. Since the packet lacks the options fields which need for DHCP, then `pfe_nhdb_dhcpv4_msg_type()` marks it as DHCPNOMSGTYPE. [PR862206](#)
- In some corner cases SPMB can get stuck in READY state. Restarting the SPMB does not help to recover from the problem state. The issue is fixed by this PR. [PR866127](#)
- mgd crashed with core file after executing "show configuration | display rfc5952". Issue is fixed now. [PR869650](#)
- On MX Series routers with DPC (I-Chip based) type FPCs running an 11.4 (or newer) Junos OS release, disabling uRPF on a logical interface might result in another logical interface on the router to drop all incoming packets. This problem happens only when the following conditions are met concurrently: a) 2 different logical interfaces share the same lookup index b) both logical interface have uRPF enabled c) these 2 different logical interfaces belong to 2 different FPCs d) at least one of the logical interfaces belongs to a DPC (ICHIP based) type FPC. The lookup index is calculated by taking the lower 16 bits of the logical interface index. In other words lookup index = IFL index MOD 65536 . It is normal, valid, and expected to have logical interfaces which share the same lookup index. The problem described in this PR is \_not\_ the fact that the lookup indexes are the same. Here is an example of 2 different logical interfaces on 2 different FPCs which share the same lookup index: Interface ge-0/1/0.945 has an logical interface index of 1774 and a lookup index 1774: `user@router-re1> show interfaces ge-0/1/0.945` Logical interface ge-0/1/0.945 (Index 1774) (SNMP ifIndex 1635) `^^^^^^^^^^ Flags: Device-Down SNMP-Traps 0x4000 VLAN-Tag [ 0x8100.945 ] Encapsulation: ENET2 Input packets : 0 Output packets: 0 Protocol inet, MTU: 4462 Flags: Sendbcst-pkt-to-re, uRPF, uRPF-loose Addresses, Flags: Dest-route-down Is-Preferred Is-Primary Destination: 52.3.168.216/29, Local: 52.3.168.217, Broadcast: 52.3.168.223 Protocol multiservice, MTU: Unlimited And interface xe-2/2/0.0 has a logical interface index of 198382 and a lookup index of 198382 MOD 65536 = 1774: user@router-re1> show interfaces xe-2/2/0.0 Logical interface xe-2/2/0.0 (Index 198382) (SNMP ifIndex 698) ^^^^^^^^^^^^^^ Flags: SNMP-Traps 0x4004000 Encapsulation: ENET2 Input packets : 381 Output packets: 376 Protocol inet, MTU: 1500 Flags: Sendbcst-pkt-to-re, uRPF, uRPF-loose Addresses, Flags: Is-Preferred Is-Primary Destination: 155.154.153.0/30, Local: 155.154.153.1, Broadcast: 155.154.153.3 Protocol multiservice, MTU: Unlimited In the example above if uRPF is disabled on ge-0/1/0.945 then xe-2/2/0.0 will start dropping all incoming packets due to RPF failure. When this condition occurs the only way to recover is to disable, commit and re-enable uRPF on the broken interface. When this is done the following error messages are generated: Apr 15 16:02:53 router-re1 fpc2 rt_iff_generic_topo_handler: jtree error Not found for disconnect on iff-post-src Apr 15 16:02:54 router-re1 fpc2 RT(rt_rpf_jtree_drt_remove_ifl): Unable to remove logical interface 198382 from drt(4) Apr 15 16:02:54 router-re1 fpc2 RT(rt_rpf_jtree_drt_remove_ifl): Unable to remove logical interface 198382 from loose(7). PR873709`
- If interface flaps of a bridge-domain with igmp-snooping are enabled or multicast snooping routes are pruned due to Designated Router changes, LUCHIP might report



traps and EDMEM read errors. These conditions are transient and only seen once the system is operating with enhanced-ip mode. [PR879158](#)

### ***Routing Protocols***

- When "passive" and "disable" knobs are both configured under **[edit protocols IS-IS interface <intf> level <N>]** hierarchy, the interface is treated as "passive" instead of being disabled. [PR697553](#)
- If you have configured PIM nonstop active routing (NSR), a core file might be created on an upstream router because of high churn in unicast routes or a continuous clearing of PIM join-distribution in the downstream router. To prevent this possibility, disable NSR for PIM. [PR707900](#)
- On a device that is running Protocol Independent Multicast (PIM) and with nonstop active routing (NSR) enabled on the device, if a PIM corresponding interface flaps continuously, a PIM thread might attempt to free a pointer that has already been freed, causing the routing protocol process (rpd) to crash and create a core file. [PR801104](#)
- Junos OS checks for mask-length mismatch for OSPF P2P-over-LAN interfaces, but skips the check if an interface has /32 mask configured. In a scenario with OSPF configured between Juniper platform and other vendor's platform, if a /32 mask IP address is configured on P2P-over-LAN OSPF interface of Juniper platform and a non /32 mask IP address is configured on the peer, the OSPF neighbor can establish but Kernel Routing Table (KRT) queue gets stuck. [PR840122](#)
- In subscriber management environment, routing protocol process (rpd) may crash and generate a core file due to snmpwalk fails at mplsL3VpnVrfRtInetCidrDestType when a subscriber access-internal route in a VRF has a datalink nexthop (such as when DHCP subscriber connects into a VRF). When issue happens, the following behaviors could be observed: user@router> show snmp mib walk ascii mplsL3VpnVrfRtInetCidr | no-more Request failed: Could not resolve 'mplsL3VpnVrfRtInetCidr' to an OID user@router> show snmp mib walk ascii mplsL3VpnVrfRtInetCidrDest | no-more Request failed: General error. [PR840323](#)
- In IS-IS scenario, with graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) enabled, after Routing Engine switchover, in very rare case, routing protocol process (rpd) might crash and generate a core file on new master (old backup) Routing Engine. This crash happens upon IS-IS lsp generation due to memory corruption. [PR841558](#)
- Under certain conditions moving a link that has BFD clients can cause stale BFD entry for the old link. [PR846981](#)
- Upstream interface of multicast rpf not matching multicast route in Inter-AS PIM. [PR847370](#)
- In multicast environment with PIM configured, in RP-on-a-stick scenario, if the rendezvous point (rp) receives multicast traffic but there are no receivers, rp's kernel will keep sending resolve requests to routing protocol process (rpd). These resolve requests might stick in resolve queue delaying other (S,G) resolves and thereby multicast traffic will be blackholed. [PR851210](#)

- If an invalid PIM-SSM multicast group is configured on the routing device, then when you issue the "commit" or "commit check" command, a routing protocol process (rpd) core file is created. There is no traffic impact because the main rpd process spawns another rpd process to parse the corresponding configuration changes, and the new rpd process crashes and creates a core file. When this problem occurs, you might see the following messages: user@router# commit check error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86) error: configuration check-out failed user@router# commit error: Check-out pass for Routing protocols process (/usr/sbin/rpd) dumped core(0x86) error: configuration check-out failed. [PR856925](#)
- RPD cored on the backup Routing Engine and BGP didn't receive complete routes. [PR863148](#)
- In Release 12.1 MPLS OAM programs BFD, it does not provide the source address(no change in behavior). In BFD before programming PPMD it queries kernel for the source address matching the prefix of the destination address on a interface. BFD programs PPMD with this source address. PPMD will construct BFD packet with BFD provided source address in the IP header. [PR870421](#)
- If a static route is configured and exported into OSPF, and it has the same subnet as an OSPF interface address, then committing configuration change (even unrelated to OSPF, like device's hostname) may cause the static route related OSPF type-5 LSA to be removed from the OSPF database. [PR875481](#)
- In inter-AS Option-B L2VPN scenario, the ASBR might create a L2VPN cloned transit route incorrectly due to a cloned route is a Juniper specific mpls.0 route which Junos OS creates on the penultimate hop router. Then in a rare case, routing protocol process (rpd) tries to delete the L2VPN cloned transit route (in mpls.0 table) multiple times. After this, routing protocol process (rpd) crashes and creates a core file. [PR878437](#)
- RPD CPU utilization keeps 100% due to "BGP resync" task when BGP is configured with no neighbor and NSR is configured. id@router> show configure routing-options nonstop-routing; id@router> show configure protocols bgp { group bgp-group { type internal; inactive: neighbor 1.0.0.1; } }. [PR884602](#)
- When used JunoScript to run command 'get-pim-neighbors-information instance=' (with NULL instance name), which triggered core file even though there are no routing-instances with pim enabled. It won't trigger core file if JunoScript command includes any instance name. [PR887070](#)

### **Services Applications**

- The jnxNatSrcNumPortInuse counter is not refreshing when polling the jnxNatSrcNumPortInuse OID via SNMP after RSP switchover. [PR829778](#)
- In L2TP subscriber management environment, after issuing CLI command "commit full", jl2tpd process (l2tp daemon) deletes all tunnel profiles and brings down all L2TP subscribers. Even though there are no configuration changes. [PR834504](#)
- MAC Flow-control asserted and MS-DPC reboot is needed. [PR835341](#)
- 1) corrected the log to state 4 bundles per tunnel to have been exhausted. 2) changed the log level from INFO to DEBUG 3) added more context to previous log: New IPSec SA install time 1356027092 is less than old IPSec SA install time 1356027092 new log

= Tunnel:<tunnel-id> <Local\_gw, Remote\_gw>: <local-gw-ip-addr, remote-gw-ip-addr> New IPSec SA install time 1356027092 is less than old IPSec SA install time 1356027092 4) added more context to previous log: SA to be deleted with index 3 is not present new log = SA to be deleted with index 3 is not present <Local\_gw, Remote\_gw>: <local-gw-ip-addr, remote-gw-ip-addr> 5) added a counter to show the number of times each of these messages occur per tunnel. [PR843172](#)

- This PR fixed syslog is not sent to remote host when rsp interface is used. [PR849995](#)
- jnxNatSrcNumSessions SNMP OID is broken in 11.4R6-S1 release. [PR851989](#)
- Defining an application with destination-port range starting at 0 can cause TCP handshake to fail through NAT. As a workaround, specify the application with destination-port range starting at 1 instead of 0. [PR854645](#)
- The number of terms per NAT rule cannot exceed 200 for the inline-service si- interface. This constraint check is not applicable for other types of service interfaces like sp-, AMS and ms- etc. Following error message will be displayed when there are more than 200 terms per NAT rule: regress@aria# commit [edit services] 'service-set ss8' NAT rule rule\_8 with more than 200 terms is disallowed for si-0/0/0.8 error: configuration check-out failed. [PR855683](#)
- Using "destination-address 0::0/0" in SFWv6 presents a commit warning. [PR857106](#)
- MS-DPC may crash in certain scenarios when using CGNAT PBA and junos-rsh, junos-rlogin, junos-rpc-services-udp and junos-rpc-services-tcp ALGs (either one) in combination with EIM. [PR862756](#)
- Service PIC might crash in corner cases when SIP ALG media flows are deleted. [PR871638](#)
- The issue is seen because of receiving malformed LCP configure-request packet with bad option length from PPP client. In this case when router tries to generate configure-nak it crashed. As a fix, check is added to discard such malformed configure-request packets. [PR872289](#)

### ***Subscriber Access Management***

- In DHCP/PPPoE subscriber management environment, after terminating subscribers, authd process might crash and generate a core file due to an invalid pointer is used. [PR821639](#)
- In situation when CoA message includes both LI attributes and CoA attributes, authd process fails to respond to CoA. [PR821876](#)
- Subnet mask option is not returned to DHCP client when framed-ip-address is used with dhcp-local-server. [PR851589](#)
- Authd core experienced when multiple DHCP subscriber connection attempts require SRC for subscriber authentication. [PR862037](#)
- Fixed the misbehavior of 'accounting-stop-on-failure' configuration knob. [PR865305](#)
- PPPoE subscribers do not always get disconnected after the client-session-timeout expires. [PR869559](#)
- DTCP - First 127 triggers are applied. [PR873013](#)

### ***User Interface and Configuration***

- On JWEB:Monitor->SystemView->System Information->General->Serial Number is missing for M10i routers. It is not displaying the serial number. On MX480 and MX240 its not displaying the proper chassis serial numbers on the above field. [PR818900](#)
- If a commit sync error occurs for a commit performed in "edit private" mode and later it is followed by another commit in global mode (without private or exclusive mode), the configuration file may remain unzipped after the global commit is complete. [PR823555](#)
- If you are using 12.2R2 or later, you might be unable to delete and commit configuration when in "configure private" mode. At that time you will see the following message. warning: patch removes statement that is not empty. [PR862959](#)

### ***VPNs***

- While l2circuit/l2vpn is not configured, if user requests for PW object info through mib, L2circuit/l2vpn is creating invalid job, which can lead to rpd crash. The fix exists in: 12.3R3, 11.4R8, 13.1R2, 12.2R5, 12.1R7 and later releases. [PR854416](#)
- When the egress PEs are on a NGMVPN, which then leads on to the assert being silently ignored, when dual forwarders are set up over the PE-CE segment. Eventually duplicate traffic being delivered by PE routers onto the ethernet where receiver is connected. [PR862586](#)
- RPD can crash when a cmcast leave is received after disabling the internet-multicast. [PR864304](#)
- Sample topology: multicast +---+ CE\_R +---+ PE\_R +---MPLS core---+ PE\_S +---+ C-BSR +---+ C-RP +---+multicast receiver source With the NG MVPN setup, when RP failed, there could be a delay on RP timeout between PE\_S (multicast traffic ingress) and PE\_R (multicast traffic egress). And suppose that PE\_S removed RP from the PR list and PE\_R still learned RP. Under the condition above, when RP came back and BSR informed RP info with generating bootstrap message, PE\_R would advertise type 6 routes to PE\_S across MPLS core via MPBGP. If a RP is learned on PE\_S after PE\_S receives the type 6 routes from the core, PE\_S neither creates PIM (\*G) join nor sends the join to C-RP. [PR866962](#)

### ***Release 12.2R4***

#### ***Class of Service (CoS)***

- When rate limit is enabled and disabled on port cos scheduler configuration leaves rate limit configuration on queues in effect. This causes the rate limit feature to be in effect even after the rate limit is removed. This PR addresses this issue in lieu of PR 843603. [PR833431: This issue has been resolved.](#)
- Traffic-control-profile-remaining is not working for logical interface in interface-set. [PR835933: This issue has been resolved.](#)
- This seems to be hard to reproduce and noticed only once after GRES. When the cosd restarts (due to the GRES test you performed), cosd reconciles the configurations pushed to the Packet Forwarding Engine with the config read from the CLI and tries to

reuse the object ID. In this case, it tries to insert the same ID twice. [PR848666: This issue has been resolved.](#)

- Commit throws an error "Invalid rewrite rule rule-name for logical interface <ifl-name> lfd <ifd-name> is not capable to rewrite inner vlan tag 802.1p bits" even though there is no rewrite configuration related to inner-vlan tag. [PR849710: This issue has been resolved.](#)

### **Forwarding and Sampling**

- There is always a chance to see this issue if any daemon adds a blob size that comes closer to 65520 (after IDR encoding). [PR700635: This issue has been resolved.](#)
- The issue of core files generated for Layer 2 Address Learning process (l2ald) is restricted to MX-FPC and might not happen in a steady condition. [PR809873: This issue has been resolved.](#)
- A memory leak might occur to the pfd, dcd, cosd, cfmd, and dfcd processes if the user frequently and repeatedly executes "show interface extensive" command from multiple telnet sessions under the following conditions: 1. Set screen-length value to small value. Screen length can be changed by the command **set cli screen-length <n>**. 2. User enters the **show interface extensive** command simultaneously from multiple telnet sessions. And cancel the output of the command with "q" as soon as **---(more)---** shows up at the end of the output. [PR843145: This issue has been resolved.](#)
- With more than four archive-sites configured under the **[edit system archival configuration archive-sites]** hierarchy level, after committing the configuration changes, the pfd process crashes and creates a core file due to memory corruption or double free. The core files could be seen by executing CLI command "show system core-dumps". [PR849465: This issue has been resolved.](#)
- MPLS forwarding table filter (ftf) is not getting linked in JTREE after router or FPC reboot. [PR851599: This issue has been resolved.](#)

### **General Routing**

- Prior to this change, the L2TP sessions with cos/ firewall attachments fail to come up when the L2TP Access Concentrator (LAC) is reachable over a unicast nexthop. [PR660208: This issue has been resolved.](#)
- A feature to support OSI/CLNS type frames on MX Series routers with MICs/MPCs was added in 11.2R1, and this feature is not supported on MX Series routers with DPC. Despite the DPC not supporting CLNS-type next hops, the following log message is seen: > \$ zcat messages.log.gz | grep 50556 > Sep 20 02:27:52 r2 fpc0 Table nexthop for unsupported proto. NH(50556) CLNP:2617 > Sep 20 02:27:52 r2 fpc1 Table nexthop for unsupported proto. NH(50556) CLNP:2617 > Sep 20 02:27:52 r2 fpc3 Table nexthop for unsupported proto. NH(50556) CLNP:2617 > Sep 20 02:53:20 r2 fpc0 NH: Failed to find nh (50556) for deletion > Sep 20 02:53:20 r2 fpc1 NH: Failed to find nh (50556) for deletion > Sep 20 02:53:20 r2 fpc3 NH: Failed to find nh (50556) for deletion This fix masks this log message for DPC. [PR680782: This issue has been resolved.](#)
- MPLS LDP traceroute does not work if you have a default route 0/0 pointing to discard on the egress router with DPC cards. [PR790935: This issue has been resolved.](#)

- On T1600-FPC4-ES, T640-FPC3, T640-FPC3-E and T640-FPC3-E2 platforms that have multiple Packet Forwarding Engines, with auto-bandwidth enabled on LSPs where CoS-based forwarding (CBF) is configured, auto-bandwidth might trigger minor changes on LSP nexthops. After this, flapping the corresponding interface or any next hop changes might result in an FPC crash and generate a core file. The core files can be seen by executing the CLI command **show system core-dumps**. This issue is seen with auto-bw configuration where there is continuous minor or major changes on LSP next hops based on traffic conditions. When this issue happens, the following logs could be seen: fpc3 PDP(pdp\_free): %PFE-3: Invalid PDP 0x4e01d7d0 fpc3 PDP(pdp\_free): %PFE-3: Error while removing PDP (0x4df4c068) fpc3 PDP(pdp\_free): %PFE-3: Error while removing PDP (0x525b3f78) fpc3 PDP(pdp\_free): %PFE-3: Invalid PDP 0x4de522b0' [PR818021: This issue has been resolved](#).
- ICMP redirects are not disabled even after configuring no-redirects on an IRB interface. [PR819722: This issue has been resolved](#).
- When an MS-DPC PIC reboots due to a crash or manual intervention, it might get stuck in a booting loop if the MS-DPC up-time is more than 49 days and 17 hours. After 5 consecutive boot failures, the MS-DPC PIC goes offline automatically and generates the following error message: [ 15:21:22.344 LOG: Err] ICHIP(0): SPI4 Training failed while waiting for PLL to get locked, ichip\_sra\_spi4\_rx\_snk\_init\_status\_clk [ 15:21:22.344 LOG: Err] CMSPC: I-Chip(0) SPI4 Rx Sink init status clock failed, cmsdpc\_spi4\_init [ 15:21:22.344 LOG: Err] CMX: I(0) ASIC SPI4 init failed [ 15:21:22.379 LOG: Err] Node for service control ifl 68, is already present [ 15:21:23.207 LOG: Err] ASERO SPI-4 XLR source core OOF did not go low in 20ms. [ 15:21:23.208 LOG: Err] ASER/XLR0 spi4 stop src train failed! [ 15:21:23.208 LOG: Err] ASERO XLR SPI-4 sink core DPA incomplete in 20ms. [ 15:21:23.208 LOG: Err] ASER/XLR0 spi4 sink core init failed! [ 15:21:24.465 LOG: Err] ICHIP(0): SPI4 Stats Unexpected 2'b 11 Error, isra\_spi4\_parse\_panic\_errors [ 15:21:24.465 LOG: Err] ICHIP(0): SPI4 Tx Lost Sync Error, isra\_spi4\_parse\_panic\_errors. To recover from this state, the whole MS-DPC needs to be rebooted. [PR828649: This issue has been resolved](#).
- PPPoE sessions cannot be established as rpd is unable to read or access profile database during access-internal route creation via "dynamic-profile->routing-instances->routing-options->access-internal" stanza. [PR830779: This issue has been resolved](#).
- When an FPC goes offline due to hardware failure and is stuck in a boot mode, it might affect Routing Engine-Packet Forwarding Engine communication on other FPCs. [PR831233: This issue has been resolved](#).
- On T4000 systems where the following conditions are met: - the [edit forwarding-options sampling input maximum-packet-length] statement is configured to a non zero value - packets are sent to be sampled from a Type 5 FPC to an ES-Type FPC housing the Multiservices PIC used for sampling, then an incorrect format of the notification header sent to the destination ES-Type FPC will trigger a packet loss in the packets sent to be sampled. The following message will be logged in the syslog on the destination FPC: [Jan 17 12:43:25.388 LOG: Err] SRCHIP(0): 1 Bad packets on p1 [Jan 17 12:43:25.389 LOG: Err] SRCHIP(0): 1 SONN errors on p1 The outcome is that the respective packets will be dropped and they will not be sampled. [PR839696: This issue has been resolved](#).

- When you configure a tunnel interface in MXVC, the tunnel interface is set to harddown. [PR839784: This issue has been resolved.](#)
- When the transit traceroute packets with ttl=1 are received on the LSI interface, you can retrieve the Source Address from the LSI interface to reply ICMP. As an LSI does not have any IFA, it uses first the IFA in the routing-instance to send a reply. Therefore, the source address used is the first IFA added in the VPN routing-instance. As a workaround, if the incoming interface is an LSI interface, retrieve the source address from the logical interface that has the destination IP Address. This ensures that the reply contains the source address from the CE facing the logical interface. [PR839920: This issue has been resolved.](#)
- Dynamic arp or routing does not work when using ether-over-atn-llc in the new PIC. [PR840159: This issue has been resolved.](#)
- Distributed protocol adjacencies (LFM/BFD/etc) might experience a delay in keepalives transmission and/or processing due to a prolonged CPU usage on the FPC microkernel on T4000 Type 5-3D FPCs. The delay in keepalive transmission/processing might result in a mis-diagnosis of a link fault by the peer devices. The issue is seen several seconds after a Routing Engine mastership switch with NSR is enabled, and the fault condition clears after a couple of minutes. [PR849148: This issue has been resolved.](#)

#### ***High Availability (HA) and Resiliency***

- The PR fix attempts to synchronize the configuration again after connection to the master succeeds if the configuration synchronization had failed earlier. [PR783832: This issue has been resolved.](#)

#### ***Infrastructure***

- The root cause of the problem was IFADDR change in VRRP context was not replicated to GRESS backup. [PR790485: This issue has been resolved.](#)

#### ***Interfaces and Chassis***

- Under certain circumstances, an MX80 device might crash when using the command "request system snapshot". [PR603468: This issue has been resolved.](#)
- vmcore is seen when file system corruption occurs during Junos OS upgrade. [PR683554: This issue has been resolved.](#)
- Kernel might cache a high incorrect value for stats and reject the correct subsequent stats coming from the PIC. The fix consists in checking if the difference of what is cached in the kernel and what is reported by the PIC is less than an acceptable value. If the answer is no, the kernel does not get stuck permanently and recovers while fetching stats subsequently. [PR806015: This issue has been resolved.](#)
- On an MX80 device, a broadcast storm on fxp0 (out-of-band interface) might cause the process "irq32:tsec2" to consume enough CPU causing the Routing Engine to lose the connection with Forwarding Engine Processor (TFEB). In such cases, the Routing Engine declares the TFEB unreachable, and the chassisd process might shutdown causing all interfaces to be removed leading to traffic loss. A few moments later, the Routing Engine is able to re-establish connectivity to the TFEB and Packet Forwarding Engine components begin to get re-initialized. However, if the broadcast storm on fxp0



still exists, this issue might occur again. If the issue occurs, the following logs could be seen: /kernel: Interrupt storm detected on "irq32: "; throttling interrupt source /kernel: peer\_inputs:3690 VKS0 closing connection peer type 17 indx 0 err 5 /kernel: pfe\_send\_failed(index 0, type 17), err=32 /kernel: pfe\_listener\_disconnect: conn dropped: listener idx=0, tnpaddr=0x80000032, reason: none chassisd[1204]: CHASSISD\_SHUTDOWN\_NOTICE: Shutdown reason: TFEB connection lost chassisd[1204]: CHASSISD\_IFDEV\_DETACH\_FPC: ifdev\_detach\_fpc(0) chassisd[1204]: CHASSISD\_SNMP\_TRAP10: SNMP trap generated: Fru Offline (jnxFruContentsIndex 7, jnxFruL1Index 1, jnxFruL2Index 0, jnxFruL3Index 0, jnxFruName FPC @ 0/\*/\*, jnxFruType 3, jnxFruSlot 0, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 0) chassisd[1204]: CHASSISD\_IFDEV\_DETACH\_FPC: ifdev\_detach\_fpc(1) chassisd[1204]: CHASSISD\_SNMP\_TRAP10: SNMP trap generated: Fru Offline (jnxFruContentsIndex 7, jnxFruL1Index 2, jnxFruL2Index 0, jnxFruL3Index 0, jnxFruName FPC @ 1/\*/\*, jnxFruType 3, jnxFruSlot 1, jnxFruOfflineReason 2, jnxFruLastPowerOff 0, jnxFruLastPowerOn 0) chassisd[1204]: CHASSISD\_IFDEV\_DETACH\_ALL\_PSEUDO: ifdev\_detach(pseudo devices: all) alarmd[1205]: shutting down chassisd connection: chassisd ipc pipe read error alarmd[1205]: chassisd alarmd[1205]: connecton succeeded after 0 retries alarmd[1205]: resending alarm state craftd[1206]: craftd\_user\_conn\_shutdown: socket 5, errno = 0 craftd[1206]: chassisd connection succeeded after 0 retries chassisd[1204]: CHASSISD\_SNMP\_TRAP7: SNMP trap generated: FRU insertion (jnxFruContentsIndex 6, jnxFruL1Index 1, jnxFruL2Index 0, jnxFruL3Index 0, jnxFruName TFEB, jnxFruType 5, jnxFruSlot 0) The cpu usage of process "irq32:tsec2" could be observed by following command: user@router> show system processes extensive | match "aver|PID|irq32" last pid: 1380; load averages: 0.76, 0.61, 0.36 up 0+00:23:47 09:46:40 PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND 31 root 1 -68 -187 OK 16K WAIT 1:03 0.63% irq32: tsec2 [PR816253: This issue has been resolved.](#)

- Prior to this PR, the speed of a GE interface capable of working at FE speeds was set to 'auto' in the Packet Forwarding Engine level. This causes a problem when manually setting the speed on the Routing Engine. Now the behavior is to set the speed to '1 g' in the Packet Forwarding Engine. For automatic speed detection the interface should be set to 'speed auto' in the configuration. [PR821512: This issue has been resolved.](#)
- The MX Series chassis-control interrupt storm might be falsely reported when a Field Replaceable Unit (FRU) is removed, inserted, or an FPM button is pushed. A FRU might not be recognized/booted, resulting in chassis operational failure. [PR823969: This issue has been resolved.](#)
- IEEE 802.3 ah LFM stats counter "OAM current frame error event information" is not cleared correctly by CLI operation. [PR827270: This issue has been resolved.](#)
- If per-unit-scheduler is configured under a physical interface(ifd), and trying to delete this ifd and its sub-interfaces (ifl) in one single commit, ksyncd might core in the backup Routing Engine which will cause GRES malfunction. [PR827772: This issue has been resolved.](#)
- If we receive a MAC Move event or an L2 logical interface change event, we don't immediately remove the next hops. The backup Routing Engine has to delete the NH first and then it gets deleted from the master. During this phase, if pointer is stale to



the NHs as in the pointer pointing to the NH is valid, but the NH has already been deleted, then you will run into this condition. [PR829093: This issue has been resolved.](#)

- A request (like snmp query) for collecting input ipv6 stats of ae logical interface on abc chipset is not working properly. [PR831811: This issue has been resolved.](#)
- Currently no SNMP trap is generated when FPC crashes. This PR is meant to enable the SNMP trap for such a failure. [PR835112: This issue has been resolved.](#)
- Currently, no SNMP trap sent when the backup SPMB fails. This PR is meant to enable the SNMP trap for such a failure. [PR835167: This issue has been resolved.](#)
- Although the physical interface is disabled, reseating 1GbE SFP on MPC/MIC restores its output optical power, hence the opposite router interface turns Up(Near-end interface is still down). Only 1g-SFP on MPC/MIC has the problem, but 1g-SFP on DPC/MX, EX Series and 10G-XFP on DPC/MX don't have the problem. When the SFP is reseated, then the SFP periodic is going ahead and enabling the laser irrespective of the fact that interface has been enabled or disabled. Driver needs to store the state for each SFP link and enable laser based on that. This software problem is fixed in Junos OS Releases 11.4R7, 12.1R6, 12.2R4, and later. [PR836604: This issue has been resolved.](#)
- Configuring 100-Gigabit Ethernet Link Down Notification for Optics Options Alarm or Warning. The "optics-options" alarm/warning "low-light"; the syslog action was not taking effect on T1600 and T4k for 100 GE PICs. This was fixed as part of this PR. [PR836709: This issue has been resolved.](#)
- It is possible that when a DPC boots up and link-training fails for all the links between fabric planes and the FPC, the interface is still brought up online and traffic blackholing will occur. [PR839076: This issue has been resolved.](#)
- The logical interfaces are marked with 0 (null) after deactivate system commit synchronize and deactivate chassis redundancy which results in the backup Routing Engine generating core files. [PR840167: This issue has been resolved.](#)
- ERA events are not credited back by jpppoed. ERA has a purge timer of 10 minutes which reclaims stale events so new connections are allowed after the purge timer fires. In a high scaled scenario this can lead to slow PPPoE connections. [PR842935: This issue has been resolved.](#)

### **Layer 2 Ethernet Services**

- When changing an interface framing from lan-phy (default) to wan-phy and back a few times, the interface might not show up any more in the **show interfaces terse** command. [PR836382: This issue has been resolved.](#)
- DHCPv6 relay terminates the client if DHCPv6-REPLY message from server contains status-code option. [PR845365: This issue has been resolved.](#)
- In certain cases when an MX Series device is configured as a DHCPv6 server and servicing DHCPv6 clients through LDRA relay, it might send advertisements with UDP port 546 instead of 547. [PR851642: This issue has been resolved.](#)

**MPLS**

- If the current configuration is - label-switched-path lsp1 { to XX.XX.XXX.XX; primary path1; secondary path2 { standby; } } And if the following configuration change is made (delete the LSP, reconfigure the LSP and make path2 as primary and path1 as standby) - delete protocol mpls label-switched-path lsp1 set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX primary path2 set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX secondary path1 standby commit It will result in a stale 'path2' standby. Later if another configuration change happens for 'path2', it can point to the stale entry and result in the assertion failure and core. The workaround is to do a 'commit' after deleting the LSP in the above configuration. Thus the configuration steps become - delete protocol mpls label-switched-path lsp1 commit set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX primary path2 set protocols mpls label-switched-path lsp1 to XX.XX.XXX.XX secondary path1 standby commit [PR847038: This issue has been resolved.](#)

**Network Management and Monitoring**

- The default maximum log file size depends on the platform type. For TX Matrix or TX Matrix Plus routers, it is expected to be 10 MB. However, due to a software defect, this file size might only be 1 MB. [PR823143: This issue has been resolved.](#)
- On a router with interfaces with Frame Relay encapsulation, an SNMP WALK operation will cause a MIB daemon (mib2d) crash and will generate a mib2d core file. The crash itself does not cause any impact on the router as the MIB daemon is restarted automatically. The only effect is that an SNMP WALK will never complete successfully. user@router-re1> show snmp mib walk 1 | no-more sysDescr.0 = Juniper Networks, Inc. mx480 internet router, kernel Junos OS 11.4R6.5 #0: 2012-11-28 21:57:12 UTC builder@evenath.juniper.net:/volume/build/junos/11.4/release/11.4R6.5/obj-i386/bsd/kernels/JUNIPER/kernel Build date: 2012-11-28 21:39:15 UTC Copyright (c sysObjectID.0 = jnxProductNameMX480 sysUpTime.0 = 339594 sysContact.0 < ..... > dot3OutPauseFrames.942 = 0 dot3OutPauseFrames.943 = 0 dot3OutPauseFrames.953 = 0 dot3OutPauseFrames.954 = 0 frDlcmilfIndex.153 = 153 frDlcmilfIndex.512 = 512 frDlcmilfIndex.513 = 513 frDlcmiState.153 = 6 Request failed: General error user@router-re1> show log messages Dec 20 09:23:20 router-re1 clear-log[8240]: logfile cleared Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3-BAD\_PAGE\_FAULT: pid 7382 (mib2d), uid 0: pc 0x810fe09 got a read fault at 0x7c, x86 fault flags = 0x4 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Trapframe Register Dump: Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: eax: 00000000 ecx: bfbeda88 edx: 00000000 ebx: bfbeda7c Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: esp: bfbeda60 ebp: bfbeda98 esi: 089de834 edi: 089fb680 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: eip: 0810fe09 eflags: 00010297 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: cs: 0033 ss: 003b ds: bfb003b es: 003b Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: fs: 003b trapno: 0000000c err: 00000004 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Page table info for PC address 0x810fe09: PDE = 0x42e60067, PTE = 5290c425 Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: Dumping 16 bytes starting at PC address 0x810fe09: Dec 20 09:23:38.683 router-re1 /kernel: %KERN-3: 8b 40 7c 89 04 24 e8 5a 3f 2f 00 89 45 ec 8b 55 Dec 20 09:23:40.787 router-re1 init: %AUTH-3: mib-process (PID 7382) terminated by signal number 11. Core dumped! Dec 20 09:23:40.787

```

router-re1 init: %AUTH-6: mib-process (PID 8247) started Dec 20 09:23:40.809
router-re1 mib2d[8247]: %DAEMON-5-LIBJSNMP_SA_IPC_REG_ROWS:
ns_subagent_register_mibs: registering 88 rows Dec 20 09:23:41.595 router-re1
mib2d[8247]: %DAEMON-6-LIBJSNMP_NS_LOG_INFO: INFO:
ns_subagent_open_session: NET-SNMP version 5.3.1 AgentX subagent connected Dec
20 09:23:43.533 router-re1 dumpd: %USER-5: Core and context for mib2d saved in
/var/tmp/mib2d.core-tarball.0.tgz Dec 20 09:23:43.793 router-re1 mib2d[8247]:
%DAEMON-6-SNMP_TRAP_LINK_UP: ifIndex 5, ifAdminStatus up(1), ifOperStatus
up(1), ifName dsc < ..... > user@router-re1> show system core-dumps
/var/crash/*core*: No such file or directory -rw----- 1 root field 680417 Dec 20 09:23
/var/tmp/mib2d.core-tarball.0.tgz /var/tmp/pics/*core*: No such file or directory
/var/crash/kernel.*: No such file or directory /tftpboot/corefiles/*core*: No such file
or directory total 1 PR835722: This issue has been resolved.

```

- Under certain conditions, duplicate SNMP indexes might be assigned to different interfaces by kernel to mib2d (Management Information Base II daemon). This might cause mib2d and other daemons such as lacpd (LACP daemon) to crash, which needs to know about the SNMP index of an interface. [PR836823: This issue has been resolved.](#)

### ***Platform and Infrastructure***

- When cscript data memory limit is exceeded when executing op, event, or commit scripts, the cscript process could reset and leave a core file when failing to allocate memory past its limits. The script would not succeed to run. In the case of a commit script the commit would not succeed. [PR722161: This issue has been resolved.](#)
- On TX/TXP multi-chassis systems, time synchronization between SCC/SFC chassis and the LCC chassis is not maintained. [PR811480: This issue has been resolved.](#)
- When using configure private with a large group definition and high number of groups, the commit process might take long to merge the configuration change with the global configuration. [PR828005: This issue has been resolved.](#)
- The **deny** commands are not working for the **show route community-name** command. [PR836624: This issue has been resolved.](#)
- This applies to all Juniper M, MX and T Series routers. In certain GRES scenarios, the backup Routing Engine might not have the complete state of the NH database from the active Routing Engine and might send duplicate NH add messages to the Packet Forwarding Engine with same NH IDs when it becomes active. This could potentially cause undesirable behavior in forwarding resulting in broken forwarding state and/or FPC core files. To limit the effect of these duplicate NH add messages, only certain duplicate NH adds messages which can be handled gracefully are allowed and all other duplicate add messages are rejected. There is no work around for this problem. [PR843907: This issue has been resolved.](#)
- On MX Series and T4000 routers, when output Filter-Based Forwarding (FBF) destined to a routing-instance is configured, the packets matched by the FBF filter might be

discarded or sent to the unintended Packet Forwarding Engine. [PR845700: This issue has been resolved.](#)

- On a device that is in the configuration private mode, when you attempt to deactivate a previously defined VLAN members list and then commit the change, the mgd process creates a core file. [PR855990: This issue has been resolved.](#)

### ***Routing Protocols***

- If maximum-paths or maximum-prefixes is configured for a route table, these limits are displayed in the output of "show route summary". In affected releases, these limits were omitted from the output of "show route summary" [PR753013: This issue has been resolved.](#)
- The routing protocol process (rpd) might crash when doing multiple GRES in combination with bgp peer flapping with large number of dampened routes. This is observed only when certain sequence criteria are met, but might not be exposed under all switchover conditions. [PR793875: This issue has been resolved.](#)
- In some specific cases, spf calculation might be incomplete because of the specific order the IS-IS LSP fragments are received. [PR797278: This issue has been resolved.](#)
- Due to duplication of the traffic, asserts will be triggered. \*G and S,G asserts are not handled properly, hence, a few assert entries will not be deleted due to the Routing Engine switchover which results in core files. ?HW type of chassis/linecard/Routing Engine. "ALL" ?Suspected software feature combination. Multicast feature ?Describe if any behavior/ change to existing function - Handle the \*G and S,G assert properly. [PR809338: This issue has been resolved.](#)
- Changes to add-path prefix-policy do not get absorbed automatically and require a manual soft-clearing of the BGP session [PR818789: This issue has been resolved.](#)
- Changing the static route configuration from next-hop to qualified-next-hop will result in the static route getting missed from the routing table. Restarting the routing process can bring back the routes but can generate a core file. [PR827727: This issue has been resolved.](#)
- Multiple route nexthops will not be returned via SNMP for ipCidrRouteTable object [PR831553: This issue has been resolved.](#)
- If LDP-SYNC <hold-down> timer is configured under IS-IS interfaces, after configuration change, the IS-IS interfaces can go to state. [PR831871: This issue has been resolved.](#)
- IPv6 address syntax on Junos OS kernel and rpd logs is violation of RFC 5952 in chapter 4.2.3 [Choice in Placement of "::"]. When the length of the consecutive 16-bit 0 fields are equal, the first sequence of zero bits MUST be shortened. For example, 2002:db8:0:0:1:0:0:1 is logged in logs as 2002:db8:0:0:1::1, but 2002:db8::1:0:0:1 is the correct representation. [PR840012: This issue has been resolved.](#)
- IS-IS reports prefix-export-limit exceeded even though the number of exported routes is smaller than the configured value of prefix-export-limit. [PR844224: This issue has been resolved.](#)
- In scenarios that use BGP to distribute traffic flow specifications, if the received flow-spec Network Layer Reachability Information (NLRI) contains an invalid argument

(such as dscp is larger than 63), routing protocol process (rpd) will generate flow-spec routes and install them in the routing table for these NLRIs; but these flow routes with invalid match conditions are rejected by the dynamic firewall daemon (dfwd) from being added to the flowspec filters. When this issue occurs, the following errors could be seen: krt\_flow\_trans\_match\_config: Failed defining match conditions 10.0.1.1,1.0.0.1,proto=6,dscp=81 krt\_flow\_trans\_term\_add: Failed adding term 10.0.1.1,1.0.0.1,proto=6,dscp=81 to filter 0x9504000 - Unknown error: 0 krt\_flow\_trans\_filter\_add: Failed sending transaction (ADD FILTER SINGLE TERM) for filter 0x9504000 \_\_flowspec\_default\_inet\_\_ to add term 10.0.1.1,1.0.0.1,proto=6,dscp=81 - Invalid argument When the BGP peer withdraws these flow routes, they will only be deleted, not freed, hence causing a memory leak. [PR845039: This issue has been resolved.](#)

- In a BGP scenario with multipath configured, if a static route which has table nexthop (such as inet.0) is configured in the same routing-instance as BGP, when an interconnect link between BGP peers is brought down or flapping, the corresponding BGP session takes 90 seconds to time out. During this period routes received over the BGP session will stay there. For a multipath transit route received from both BGP sessions, initially both paths are resolved over the interconnect links directly. When one of the interconnect links is brought down or flapping, that path will be resolved over the static default route which has table nexthop (such as inet.0). Therefore, one path is resolved over a router nexthop and the other path is resolved over a table nexthop. This will cause the routing protocol process (rpd) crash and generate a core file. This issue usually occurs in BGP/L3VPN environment. The core files could be seen by executing the **show system core-dumps** command. [PR851807: This issue has been resolved.](#)

### Services Applications

- SIP ALG was not allowing SIP 603 decline message. [PR822679: This issue has been resolved.](#)
- j12tpd crash \_thr\_send\_sig (thread=0x8a5e000, sig=6) at `../../../../src/bsd/lib/libthr/thread/thr_kern.c:91` j12tpd crash exhibited in environment where MX480 was configured as LAC and terminating 500 I2tp subscribers. [PR824760: This issue has been resolved.](#)
- Making SDG1 from standby to master with all the flows already in sync might result in the CLI hang for sometime until the flows get cleared and resynchronized in SDG1 during this CLI hang timeframe. [PR829950: This issue has been resolved.](#)
- With RTSP ALG enabled, RTSP keep-alive packets might be dropped if it's already Ack'ed by the receiver. [PR834198: This issue has been resolved.](#)
- In a Carrier Grade NAT (CGNAT) scenario, without any configuration change, under some conditions, MS-DPC PIC might crash and generate a core file when encountering unknown flow-type. Service will be impacted during the period. When this issue happens, the following logs could be seen: chassisd[1477]: CHASSISD\_SNMP\_TRAP10: SNMP trap generated: FRU power off (jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index 2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @ 5/1/\*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 8, jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 33404122) chassisd[1477]: CHASSISD\_SNMP\_TRAP10: SNMP trap generated: FRU power on (jnxFruContentsIndex 8, jnxFruL1Index 6, jnxFruL2Index

2, jnxFruL3Index 0, jnxFruName PIC: MS-DPC PIC @ 5/1/\*, jnxFruType 11, jnxFruSlot 5, jnxFruOfflineReason 2, jnxFruLastPowerOff 192338801, jnxFruLastPowerOn 192338924) [PR834899: This issue has been resolved.](#)

- In scenarios that use the sp interface, such as IPsec VPN, multiservice process (mspd) will memory leak during sp interface flapping. The memory usage of the mspd process can be checked by the following CLI command: `user@router> show system processes extensive | match "PID | mspd"` (Note: The "RES" field means "Current amount of resident memory, in kilobytes")  
PID USERNAME THR PRI NICE SIZE RES STATE TIME  
WCPU COMMAND 2048 root 1 96 0 36216K 34820K select 0:10 0.00% mspd  
When the memory usage of the mspd process increases to the system limit (about 131072KB), the following logs could be seen: `/kernel: %KERN-5: Process (2048,mspd) attempted to exceed RLIMIT_DATA: attempted 131076 KB Max 131072 KB` [PR836735: This issue has been resolved.](#)
- The "hot-standby" CLI knob under [**edit interfaces <RSP interface name> redundancy-options**] is made hidden for the Redundant Service PIC (RSP). [PR838762: This issue has been resolved.](#)
- Service PIC might crash under certain race conditions when receiving sip invite packets. [PR843047: This issue has been resolved.](#)
- Service PIC might crash in corner cases when receiving specific SIP REGISTER. [PR843479: This issue has been resolved.](#)
- Service PIC might crash in corner cases when EIM is enabled for SIP ALG. [PR847124: This issue has been resolved.](#)
- When allocating the memory from shared memory for bitmaps used in port blocks, the Junos OS requests as many bytes as the size of the block. For example, if you assign 10K bytes block size for deterministic NAT or PBA, then the Junos OS allocates 10K bytes for that bitmap. However, it only needs 10K/8 bytes as one byte can represent 8 ports. These huge allocations are leading to memory depletion when many source addresses are behind the NAT, and port blocks are big. [PR851724: This issue has been resolved.](#)
- The spd process generates a core file during switchover if a CGNAT is configured. [PR854206: This issue has been resolved.](#)

### ***User Interface and Configuration***

- The output of the **show system users no-resolve** command displays the resolved hostname. [PR672599: This issue has been resolved.](#)
- On J-Web, Monitor > SystemView > System Information > General > Serial Number, the serial number is missing for M10i routers. Additionally, for MX480 and MX240 routers, this field does not display the proper chassis serial numbers. [PR818900: This issue has been resolved.](#)

### ***VPNs***

- In BGP-MVPN, when the number of multicast routes falls below the threshold, the earlier suppressed MVPN multicast routes because of limit are not added back again. For MVPN, there was no mechanism to trigger the processing of cmcast entries that

were not added earlier. The fix is to queue the cmcast entries that are suppressed for multicast route addition in a new list. When the reuse limit is reached, this list is walked and used to add back the entries. [PR841105: This issue has been resolved.](#)

- In l2circuit (Martini l2vpn) scenarios where a backup neighbor is being defined along the 'standby' knob, after deleting this backup neighbor from the configuration, its associated vc-route is not being eliminated. Later if the user deletes the l2circuit neighbor or restarts the routing protocol process (rpd), the rpd process might crash and generate a core file. [PR841522: This issue has been resolved.](#)
- Deleted logical interfaces might not be freed due to references in MVPN. [PR851265: This issue has been resolved.](#)

### **Release 12.2R3**

#### **Forwarding and Sampling**

- The **show interface aeXXX detail** command might not have the correct link information and some statistics might be shown as 0 when there is traffic. [[PR828155: This issue has been resolved.](#)]

#### **General Routing**

- cgatool should not be part of jkernel-common.manifest.in Whatever product that requires cgatool would need to package jcrypto.manifest. Currently, jkernel-common.manifest.in is used by both jkernel-qfx and jkernel-ppc. jkernel-qfx does not require cgatool to be in jkernel-common.manifest.in. Need same confirmation from jkernel-ppc.manifest. [[PR549623: This issue has been resolved.](#)]
- In scaled VPLS scenario, with GRES and NSR enabled, with end-to-end traffic running, as part of the Routing Engine switchover on a CE/PE router, when the Packet Forwarding Engine receives a request from the Routing Engine requesting whether the group of MAC addresses should be aged and if the route for one of the MAC addresses from the list (not last) is not found, then the Packet Forwarding Engine replies to the Routing Engine with the subset of MAC addresses to be aged containing zero MAC address (for which route was not found) which can lead to a vm core file, causing the router to go to db prompt. [[PR783099: This issue has been resolved.](#)]
- After the Routing Engine switchover or a reboot, MLFR interfaces might not send MVPN data to all downstream receivers [[PR787168: This issue has been resolved.](#)]
- On an MX960 router with redundant Routing Engines, when you swap the members of a link aggregation group (LAG), a vmcore or ksyncd core file might be created on the backup Routing Engine. [[PR793778: This issue has been resolved.](#)]
- After a software upgrade to the Junos OS 11.4R3 or later, an xe interface might report **L2 channel errors** when the **show interface extensive** command is issued. This issue has been noticed when the xe interface is part of an "ae" bundle. [[PR800634: This issue has been resolved.](#)]
- Transient SL chip errors might be encountered when a link flap event occurs. [[PR812092: This issue has been resolved.](#)]



- Framing mode change on 8xOC3/2xOC12 ATM MIC does not trigger an automatic MIC bounce. This causes traffic issues after mode change. [[PR814856](#): This issue has been resolved.]
- When migrating from 12.2 to 12.3 using unified ISSU, the blobs being created in 12.2 are using a newer format which is not compatible with 12.3. [[PR818947](#): This issue has been resolved.]
- When an aggregated Ethernet (ae) interface is configured as an untagged bridge interface (family bridge or vpls) and when a member port of ae is disabled and enabled, the enabled member port will discard an LACP control packet. As a result, MUX status of the member port stays at DETACHED and it does not join the aggregated Ethernet. [[PR825312](#): This issue has been resolved.]
- Sometimes the new 100GE interface might be unable to display the accurate input PPS value. [[PR826596](#): This issue has been resolved.]
- VC-Subscriber license change not updated to all the Routing Engines in backup members [[PR835039](#): This issue has been resolved.]
- When you configure a tunnel interface in MXVC, the tunnel interface is set to harddown. [PR839784](#): This issue has been resolved.
- Dynamic arp or routing does not work when using ether-over-atn-llc in the new PIC. [[PR840159](#): This issue has been resolved.]

### **Infrastructure**

- If running in a low memory environment, while receiving IP fragmented packets destined to the Routing Engine, the system might fail to defragment the packets and might free the memory of these fragmented packets. But the router might try to free the memory again, and in rare cases, the freed memory might be allocated before the double free process occurs. In such scenarios, an access to the freed memory might cause memory corruption and a kernel crash. [[PR810434](#): This issue has been resolved.]
- In an L3VPN environment, when PE enabled with composite-nexthop receives an ICMP packet from local CE having TTL=1 and "route-record" option destined towards remote VPN end, then the following messages are seen in the syslog of the PE router. /kernel: %KERN-3: tag\_send\_nh\_chain(): no mbuf after tag\_nh\_chain\_comp\_label\_output() This syslog message does not indicate a real mbuf issue. Hence, this can be treated as non-intrusive. [[PR811406](#): This issue has been resolved.]
- A kernel crash might occur on routers running the Junos OS 10.4 or higher (which does not have fix for this PR), with "targeted-broadcast" knob configured on a broadcast interface. If this knob is configured, MAC address will be learned for subnet broadcast IP (configured on that interface). When this ARP table entry gets timed out, it corrupts an internal data structure, leading to kernel crash. This MAC learning will happen with one of the following : 1. Mismatched IP subnet is configured on one of the connected devices 2. A malformed packet (ARP request to subnet broadcast IP) is received on that interface NOTE: - MAC address learned for the subnet broadcast IP cannot be seen using the **show arp** command. - This issue is platform independent. [[PR814507](#): This issue has been resolved.]



- Over time when routing-table instances are added and removed, when the global index allocated to a route-table reaches around 36730, the default route-table (inet.0), having index 0 always, could be incorrectly allocated to a newly added instance. After this corruption to the default routing-instance whenever this default route-table (inet.0) is accessed or modified, this could result in kernel panic. The fix addresses this corruption by adding the default route-table inet.0 (index 0) to the structure tracking all default tables in the system, thereby ensuring that this index never gets re-allocated and does not result in this kernel core file. The possibility of running into this issue depends solely on the rate of adding and removing routing-instances. [[PR829412](#): This issue has been resolved.]
- Logical interface inet6 protocol might be stuck at down state because of either external loopback or duplicate inet6 address detected. DAD will not run after this inet6 protocol-down event. [[PR834027](#): This issue has been resolved.]

### *Interfaces and Chassis*

- After an MX Series Routing Engine reboot, the internal em interface might be "link down". The fix for this is to monitor the link state of the em interfaces. When a em link is discovered to be harddown, re-initialize the em interface. This fix has been limited to address only MX Series platforms. [[PR611081](#): This issue has been resolved.]
- On an E1 interface, when the interface flaps on the CE side of the connection, the interface will flap a second time on the PE side. [[PR690403](#): This issue has been resolved.]
- Multiple inbound/outbound IPsec tunnels are seen for a single security association upon certain conditions such as router reboot. [[PR730174](#): This issue has been resolved.]
- Certain changes to NAT (Network Address Translation) PBA (Port Block Allocation) configuration require a reboot of the services PIC in order for the changes to take effect on the PIC. A warning should be issued during the commit, informing you of the need to reboot.

For the PBA or Deterministic NAT configuration changes to take effect on the service PIC, you must reboot the services PIC when the following NAT config changes are made under the **[edit services nat pool]** hierarchy level.

- active-block-timeout
- address/address-range
- block-size
- max-blocks-per-address
- port range

[[PR807350](#): This issue has been resolved.]

- When a PEM module is inserted into a chassis, no message is currently logged in the /var/log/inventory file. This PR addresses this requirement. On inserting a PEM module into a chassis a message in the following format will be logged in file /var/log/inventory : <Current-Date> PEM <pem-slot-number> - part number <part-number>, serial

number <serial-number> eg. Oct 22 14:44:21 PEM 0 - part number 740-123456, serial number VK11111 [[PR808450](#): This issue has been resolved.]

- "hold-time down 300" does not work when alarms appear below 200ms in 1xOC192/STM64 PICs [[PR815464](#): This issue has been resolved.]
- "show interfaces redundancy" might display secondary as down upon the following sequence: deactivate R.I.(that contains entire mfr logical interfaces)-->restart fpc(that holds secondary MS pic)--> activate the R.I. back [[PR816595](#): This issue has been resolved.]
- If LACP is configured in distributed mode, which is the default mode, LACP packets are not counted in input statistics of interface [[PR821874](#): This issue has been resolved.]
- With this fix, IPv6 VRRP will not inter-op across new releases and old releases "with version 3 disabled", because of a correction in checksum calculation. [[PR826734](#): This issue has been resolved.]

### **Layer 2 Features**

- The issue is due to another PR-686399, where the DA mac entries on AE are getting aged out though the traffic is ingressing on different Packet Forwarding Engines of same AE. [[PR802924](#): This issue has been resolved.]
- In a Layer 2 VPN scenario, if there is no mpls route to neighbor and there is a static route with discard nexthop in inet.3 table as follows: user@router# show routing-options rib inet.3 { static { route 0.0.0.0/0 discard; } } Then the L2vpn connection will use the above static route in inet.3 table to connect its neighbor as follows: user@router# run show route table mpls mpls.0: 7 destinations, 7 routes (7 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, \* = Both 0 \*[MPLS/0] 00:01:42, metric 1 Receive 1 \*[MPLS/0] 00:01:42, metric 1 Receive 2 \*[MPLS/0] 00:01:42, metric 1 Receive 13 \*[MPLS/0] 00:01:42, metric 1 Receive ge-0/0/1.601 \*[L2VPN/7] 00:01:38, metric2 0 Discard

In this situation, the routing protocol process (rpd) generates a core file while walking snmp mib 'jnxVpnPwEntry'. [[PR816821](#): This issue has been resolved.]

### **Layer 2 Ethernet Services**

- In an MX Series Virtual Chassis environment, a traffic outage longer than 2 minutes occurs when a member of the VC-M LAG fails while LACP is in active mode with link protection on the VC-M router. The outage occurs while the LACP process restarts on the new VC-M router. To avoid this situation, make sure that LACP is running in active link protection mode on the device in front of the VC-M router. This device cannot be an EX Series Ethernet Switch, because the switch does not support LACP in link protection mode. [[PR784965](#): This issue has been resolved.]
- jdhcp daemon will crash which is caused by dereference of NULL client parameter. [[PR832125](#): This issue has been resolved.]
- In a scenario where a DHCP stray request is received on MX Series router acting as DHCP relay and authentication for this request fails, MX Series is generating DHCP NACK back to DHCP client [[PR835794](#): This issue has been resolved.]

**Multiprotocol Label Switching (MPLS)**

- Configuration changes to some attributes of the standby secondary path of an MPLS label-switched path (LSP) might cause the LSP to flap, which might result in packet loss. [[PR394184](#): This issue has been resolved.]
- If ldp-tunneling is configured on an LSP, and the router receives thousands of LDP routes through this LSP, and additionally the router has a large number of FPCs (especially in multi-chassis platform), during MPLS statistics collection, it is possible to see packets drop or delay on the link connecting the Routing Engine with Packet Forwarding Engine. If auto-bandwidth is enabled for the LSP, then auto-bandwidth will work incorrectly. The following kernel message type could be seen when issue occurred: /kernel: rpd pid 16252 tid 100130 syscall 133 ran for 568 ms Also, too much kernel resource was occupied once the issue occurred. The Routing Engine-based keep-alive processing could be impacted over the syscall period. bfd[4325]: BFDD\_TRAP\_MHOP\_STATE\_DOWN: local discriminator: 1715, new state: down, peer addr: x.x.x.x [[PR785360](#): This issue has been resolved.]
- When an AE link along a path becomes oversubscribed due to failure of one or more member links of the AE link, the ingress of a LSP may continue to use the same path although there might be an alternate path available with sufficient bandwidth. This is so because CSPF algorithm during optimization of an existing LSP will continue to see such an over-subscribed link acceptable with sufficient available bandwidth. However, if a new LSP with a bandwidth requirement is signaled over such a link, the LSP will not get signaled successfully. [[PR807670](#): This issue has been resolved.]
- RSVP interface state will show as DOWN on unnumbered p2p interfaces. This will cause MPLS LSPs over those interfaces to stay down. [[PR814071](#): This issue has been resolved.]
- BFD session between PEs might not come up correctly after ppmmd crashes. [[PR826300](#): This issue has been resolved.]

**Network Management and Monitoring**

- The issue has been fixed where incorrect query can cause the Routing Engine CPU utilization to go high. [[PR771867](#): This issue has been resolved.]
- There are compilation problems with the following 3 MIBs: 1. mib-jnx-license 2. mib-jnx-sp-nat 3. mib-jnx-subscriber In 11.2 version of the JUNIPER-SMI, these three objects are defined, but they are missing in 12.1. Issue has been resolved in later releases. [[PR794327](#): This issue has been resolved.]
- The 'timestamp year msec' command in syslog (not using structured data) is intended to include year and msec details in local syslog messages stored in rotating files, but not to be included in messages sent to a remote syslog collector. This fix corrects an incorrect behavior introduced in 11.3R1 where such details were also included in syslog messages sent to a remote host. [[PR820436](#): This issue has been resolved.]
- Expand the buffer size and set break point to allow sending out large snmp messages due to ospf down event. [[PR827660](#): This issue has been resolved.]

### ***Platform and Infrastructure***

- MPC might crash during unified ISSU. [[PR792909](#): This issue has been resolved.]
- The output of the following commands: - "cli > show route forwarding-table vpn <vpn-name> interface-name <interface-name>" - rtinfo might show negative values for Ipkts and Opkts counters due to invalid format of output for the variables. For example, user@test> show route forwarding-table vpn test-vpn interface-name ge-1/0/8 Name Mtu Network Address Ipkts Ierr Opkts Oerr Coll ge-1/0/8 1522 <Link> 00.1d.b5.27.48.ad -891806008 0 -1176087381 0 0 The fix also corrects the output format for input and output bytes fields. [[PR798999](#): This issue has been resolved.]
- CLI cosmetic errors are noticed while executing commands under "edit system class login" [[PR812022](#): This issue has been resolved.]
- HTTP-get probes fail when routing-instance is used for the probes. Removing routing-instance should work. [[PR814357](#): This issue has been resolved.]
- When two irb interfaces with the same Layer 2 trunk interface within the bridge domain, multicast replication might be handled incorrectly. [[PR823435](#): This issue has been resolved.]
- The timestamp in the DDOS\_VIOLATION report is not correct. The time the violation occurred is the time when the message appeared in the log. [[PR828085](#): This issue has been resolved.]
- When a MX Series based card is receiving an IPv6 packet encapsulated in an MPLS stack formed by two explicit-null labels (0 ? Ipv4 transport label, 2 ? ipv6 explicit-null label) is corrupting it. [[PR830209](#): This issue has been resolved.]
- With DHCP/BOOTP relay agent configured under the [forwarding-options helpers] hierarchy level, interface flapping will cause forwarding UDP daemon (fud) memory leak. The memory usage of the fud process could be seen by following command: (SIZE: Total memory size of the process (text, data, and stack), in kilobytes)  
user@router> show system processes extensive | match "pid | fud" PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND 8436 root 2 0 2668K 1708K select 0:03 0.00% 0.00% fud user@router> show system processes extensive | match "pid | fud" PID USERNAME PRI NICE SIZE RES STATE TIME WCPU CPU COMMAND 8436 root 2 0 2676K 1716K select 0:03 0.00% 0.00% fud [[PR831965](#): This issue has been resolved.]
- NPC core file generated at  
ns16550\_write,system\_console\_nputs,console\_putc\_polled,print\_string. [[PR835759](#): This issue has been resolved.]

### ***Routing Protocols***

- In a scaled RIP topology (about 50 RIP neighbors and route injection), with nonstop active routing (NSR) enabled, after graceful Routing Engine switchover (GRES), RIP packets will not be sent out for a few seconds to a few minutes. The blackout duration would vary by number of logical interfaces or routes the router has. The following errors could be seen when this issue happens: Cannot send request for whole table to neighbor ae1.2460: Socket is not connected send msg failed: Socket is not connected [[PR754140](#): This issue has been resolved.]
- In a scaled setup, with many routing-instances, if the FPCs are restarted, some pim encapsulation interfaces (pe) will be marked as down. [[PR770213](#): This issue has been resolved.]
- The routing protocol process (rpd) generates a core file during commit where NSR and GRES are disabled but backup rpd is still running. [[PR794738](#): This issue has been resolved.]
- In some specific cases spf calculation may be incomplete because of the specific order the IS-IS LSP fragments are received. [[PR797278](#): This issue has been resolved.]
- Enabling the "advertise-external" knob towards the Route-Reflectors or remote PEs can trigger RPD to core if: - "multipath vpn-unequal-cost equal-external-internal" enabled in a VRF - a multipath route exists with contributors from IBGP - an external path with lower preference (local-preference, as-path, etc) exists in the same VRF [[PR823844](#): This issue has been resolved.]
- During unified ISSU, LACP timeout may occur when time between fpc-upgrade and the Routing Engine switchover takes more than 90sec due to large scale config. [[PR826984](#): This issue has been resolved.]
- Any protocol (MVPN,PIM, and so on) that does install multicast next hops and goes through a forwarding-table export policy with an install-nexthop policy action will fail to install the forwarding nexthop and multicast traffic is not forwarded. [[PR830448](#): This issue has been resolved.]

### ***Services Applications***

- An MS-DPC sending DPD triggers for tunnels that does not have IPsec SAs. This results in the kmd process on the Routing Engine continuously trying to initiate an IPsec phase 1 negotiation since there are no active IKE/IPsec SAs. [[PR754461](#): This issue has been resolved.]
- Called-station-id always sent by LNS irrespective of configuration on LNS. The knob of "excluding" this attribute did not work. The fix was tested working in 11.4X27.38 [[PR818899](#): This issue has been resolved.]
- An LNS configuration on MX Series routers can be reported as invalid even if it has been successfully committed previously. [[PR823709](#): This issue has been resolved.]
- On an MX480 router, the TCP ALGs such as FTP/PPTP operations fail on a services PIC/DPC. This caused by the services PIC/DPC failing to discard expired TCP flows, causing resources exhaustion which causes failures to create additional NAT flows. [[PR825355](#): This issue has been resolved.]

- When an MX Series router uses MS-DPC to provide the tunneling service for flow-tap traffic, if there is SCU/DCU configured on the same slot of the flow-tap traffic ingress interface, all the flow-tapped sampled packets will be dropped. It is caused by the wrong nexthop linking when DPCU is configured. [[PR825958](#): This issue has been resolved.]
- Service PIC might crash in routing loops scenarios because of SIP ALG. [[PR830070](#): This issue has been resolved.]
- The lawful intercept message is cosmetic and does not impact functionality. [[PR830457](#): This issue has been resolved.]
- In the case of a stateful proxy, two SIP users behind the NAT device (so-called SIP hairpinning) will be unable to signal the call. [[PR832364](#): This issue has been resolved.]
- In the case of a transparent proxy, two SIP users behind the NAT device (so-called SIP hairpinning) will be able to signal the call, but the RTP voice flow *may* be unidirectional if one side started to send his RTP traffic before the port was opened on the other end, causing an ICMP unreachable which confuses the NAT device. [[PR834933](#): This issue has been resolved.]

#### ***Software Installation and Upgrade***

- While performing a unified in service software upgrade (ISSU), the following error appears: "[ : -a: unexpected operator]." It does not impact any functionality and can be safely ignored. [[PR806399](#): This issue has been resolved.]

#### ***Subscriber Access Management***

- Snmpwalk requests sent to MX Series routers return multiple duplicate records for jnxUserAAAAccessPool. [[PR840640](#): This issue has been resolved.]

#### ***User Interface and Configuration***

- When you configure deterministic port block allocation preserve-range settings and traffic is first sent for preserve-range ports (1023 or below) and then the non-privileged regular ports (1024-65535), packets are discarded at slow path and drop flows are created. [[PR795609](#): This issue has been resolved.]
- On J-Web:Monitor->SystemView->System Information->General->Serial Number is missing for M10i routers. It's not displaying the serial number. On MX480 and MX240 it's not displaying the proper chassis serial numbers on the above field. [[PR818900](#): This issue has been resolved.]
- 1. On the CLI prompt, hold the meta key (ALT in case you are a Windows to putty to the server) and keep pressing 5 for 30 seconds (both ALT and 5 pressed for 30 seconds)  
2. Release both keys 3. Now press 4 4. CLI would generate a core file. [[PR828046](#): This issue has been resolved.]

#### ***VPNs***

- Junos OS 11.4x26.1 upgrade might fail if the PIM MVPN instance does not have the vpn-group-address configured. [[PR753863](#): This issue has been resolved.]

- With "vpn-apply-export", per-nexthop label, and a "bgp policy overridden nexthop rather than the bgp local-address", labeled route for L3VPN did not install in standby the Routing Engine even with NSR configured. BGP on the standby Routing Engine thinks it is not the nexthop and does not create a labeled route on standby the Routing Engine. This will cause a label/nexthop change after switchover resulting in traffic loss. [[PR807285](#): This issue has been resolved.]
- For a static pseudowire, if the user configures "static send-oam", the pseudowire goes down (state: Vc-Down) [[PR829666](#): This issue has been resolved.]

## Release 12.2R2

### Class of Service (CoS)

- A memory leak in cosd can occur due to one of the cosd 8011p rewrite functions not releasing memory after use. This results in COSD memory running out of memory space thereby resulting in COSD core files continuously. [[PR782728](#): This issue has been resolved.]
- During addition/deletion or just deletion of interface configured for shared scheduler, some portion of memory is not reclaimed back. Continuous addition/deletion of these interfaces results in memory depletion causing packet loss and other issues. [[PR803939](#): This issue has been resolved.]
- On IQ2E PIC, if the logical interfaces under two different interface-sets have their own scheduler-map configured, when the same iflset shaping parameters are applied on these two interface-sets, the iflset shaping profile is not shared between the interface-sets, hence, the number of iflset supported is reduced. There is no workaround. [[PR804158](#): This issue has been resolved.]
- If you set following, Cosd crashes after executing **show class-of-service scheduler-map**. Example:  

```
*****
**** set class-of-service interfaces ge-2/0/* scheduler-map-chassis derived set
class-of-service interfaces ge-2/0/1 scheduler-map-chassis p0 set class-of-service
interfaces ge-2/1/* scheduler-map-chassis derived set class-of-service interfaces
ge-2/1/1 scheduler-map-chassis p0
*****
**** Scheduler-map-chassis for ge-2/0/1 and ge-2/1/1 has to be overrphysical
interfaceen because of the more specific configuration. This is happening correctly,
but the 'derived' scheduler-map previously allocated for these interfaces is not getting
freed. Cosd crashes when trying to display these scheduler maps since they are not
fully populated. [PR807593: This issue has been resolved.]
```
- When configured **scheduler-map-chassis derived** and auto bandwidth computation is not successful, cosd might crash after the **show class-of-service scheduler-map** command. [[PR811586](#): This issue has been resolved.]

### Forwarding and Sampling

- In a heavily scaled setup when dfwd is busy processing the filter configuration, pppmd daemon would wait (approximately 2 minutes) for firewall daemon to process the

message it sent. The wait will happen for the nth message sent, since there is a limited buffer between the two daemons. [[PR769452](#): This issue has been resolved.]

- DCU statistics can be broken on a sub-interface under certain circumstances when another sub-interface having DCU enabled gets modified (add/delete/change). By DCU statistics broken it could mean the value gets corrupted and continues to count from there properly or always remain broken. The circumstances when this issue could be triggered are: - when 2 sub-interfaces have DCU enabled - when the last 16 bits of the 20 bit sub-interface index are the same - configuration change is made on one of these sub-interfaces - other sub-interface could have DCU statistics broken. keyword: sub-interface = logical interface = IFL. When this happens we can see messages like the following being logged on the FPC syslog or messages log file. Aug 6 11:12:32.649 faraday-re1 fpc4 RT(rt\_dest\_class\_stats\_read): unable to read stats (ifl 196619, IPv4) - Not found Aug 6 11:12:32.723 faraday-re1 fpc1 RT(rt\_dest\_class\_stats\_read): unable to read stats (ifl 196619, IPv4) - Not found Aug 6 11:12:32.768 faraday-re1 fpc4 RT(rt\_dest\_class\_stats\_read): unable to read stats (ifl 196619, IPv4) - Not found Aug 6 11:12:32.843 faraday-re1 fpc1 RT(rt\_dest\_class\_stats\_read): unable to read stats (ifl 196619, IPv4) - Not found [[PR794116](#): This issue has been resolved.]
- On an MPC sometimes the policer configuration does not get correctly programmed in Packet Forwarding Engine. The issue can affect MX Series routers with MICs/MIPs with Junos OS Release 11.1 and later. It can be triggered by the following steps: 1) An interface-specific filter that contains at least a policer is already applied on an interface. 2) The parameters of the policer that is included by this filter are changed. Note that the policer type change will not trigger this issue, e.g., from term-specific policer to filter-specific policer. 3) In the instances of the interface-specific filter created after policer parameter change, the policer might not work as expected. Workaround: after 2) and before 3), make any change on the interface-specific filters that contain the changed policer. [[PR819465](#): This issue has been resolved.]

### **General Routing**

- The output of the command 'show lldp neighbors' displays contents of Port-Description TLV instead of displaying contents of mandatory PortId TLV. Also the Port Description TLV gets generated using Port Name instead of using configured Port Description. [[PR550544](#): This issue has been resolved.]
- With the configuration of VPLS interface, the NH topology has a default filter-class being associated with the logical interface in the VPLS next-hop topology chain. During the PIC Offline/Interface deactivate event, the first message received from the kernel by Packet Forwarding Engine is to delete the VPLS family from the interface. This triggers a topology change wherein the filter-class gets deleted from the topology tree. During this topology change the key buffer pointer adjustment fails thereby causing the increment in truncated key error counters for the Packet Forwarding Engine. The increment in these error counters triggers a chassis alarm associated with the R-chip error counters. These error counters stop incrementing once the interface completely goes down as a part of the PIC Offline/Interface deactivation event. [[PR718591](#): This issue has been resolved.]
- When large CLI output is displayed the CLI CPU may go high and remain high. The issue has been fixed in latest releases. [[PR731647](#): This issue has been resolved.]



- The crash is triggered by an AE link flap. The pre-requisite for the crash is a forwarding topology with composite nexthop pointing to unilist pointing to an aggregated interface (e.g VPLS pseudowire configured on MPLS LSP with node-link protection enabled). [[PR746509](#): This issue has been resolved.]
- When configuring FIB localization, if a Packet Forwarding Engine is configured with FIB-remote, heap memory leak occurs on the Packet Forwarding Engine while installing and removing some prefixes using no-route-localize policy. [[PR756787](#): This issue has been resolved.]
- On the M120, Fatal HSL2 errors might be seen on RFEB, this issue is especially seen after few days later if switchover FEB-FPC redundancy. [[PR770326](#): This issue has been resolved.]
- DPC might randomly crash during unified ISSU. It will be kept offline after the unified ISSU period. [[PR773960](#): This issue has been resolved.]
- When a large number of subscribers log in and log out simultaneously in a scaled configuration, errors might occur when logical interfaces are created. The errors occur because the rpd process fails to read ifstate notifications related to logical interface deletions. As a workaround, restart the rpd. [[PR775033](#): This issue has been resolved.]
- Routing protocol process (rpd) might crash with core file after executing command **show route community-name** with an empty string: **show route community-name**. [[PR776542](#): This issue has been resolved.]
- On T Series Core Routers and on TX Matrix and TX Matrix Plus routers, packets marked with the error flag, example due to DA (Destination Address) reject, are also counted as L3 incomplete, which is not correct and is misleading. Packets marked as error from the PIC are now counted via the 'show pfe statistics error' command. [[PR782070](#): This issue has been resolved.]
- AE member link is rejecting packets due to bad unicast MAC [[PR783332](#): This issue has been resolved.]
- On LMNR and FPCs, when a transit packet's (300 Bytes or more packet size) ingress and egress interfaces are in the same Packet Forwarding Engine (with scaled egress NHs) and if a notification is sent to the Routing Engine for that packet, FPC might reset. Following list provides some scenarios, where a notification is sent to the Routing Engine: 1. IP options packet is received 2. TTL expired packet is received 3. Sampling is configured and a packet is sampled [[PR785143](#): This issue has been resolved.]
- In scenario with indirect nexthop used (such as BGP, L3VPN), under high memory usage, in some conditions (such as interface or protocol flapping), the target nexthop of indirect nexthop changes between unilist (for example, over AE interface, or need to load balance) and unicast or just between two unilists, which have different nexthop count, Packet Forwarding Engine might generate multiple core files due to memory leak or memory corruption. [[PR787570](#): This issue has been resolved.]
- Junos OS Release 11.4 introduced PTP feature and this in turn introduced Periodic reading of clocking chip registers over SPI interface. This is a known thing since day 1 of 11.4. With newer chassis (MX80/midRangius) program (MX80-T, MX40-T, MX10-T, MX5-T) hardware is reworked to implement SPI reads in hardware and it has

significantly reduced clksyncd CPU usage to around 1-2% of CPU. [[PR789804](#): This issue has been resolved.]

- In Junos OS Release 11.4R4, ETH-DM packets greater than 994 bytes fail DMM test. The size of default ETH-DM packets is much smaller, and bigger ETH-DM packets are used only with optional data payload size. For packets sizes less than 994, the functionality works fine. [[PR790040](#): This issue has been resolved.]
- On an MX Series router, error occurs while deleting "protect protocols l2circuit" from a Virtual Chassis configuration. [[PR794782](#): This issue has been resolved.]
- This issue depends on FPC type. There are two types of FPC supported in a T4000 router: Enhanced Scaling FPC types, which are designated by the "-ES" suffix in their description and the T4000 FPC5 type FPC, which is designated by the "-3D" suffix in their description. The issue manifests in two different ways: a) Multicast traffic entering the router on an "-3D" FPC and leaving the router on an "-ES" FPC, will experience packet loss at specific packet sizes. b) If output sampling is performed on a services PIC located on an "-ES" FPC, where the ingress and egress FPCs of the traffic being sampled are distinct "-3D" FPCs, these samples can be discarded at specific packet sizes. This means flow collectors will register fewer flow records than expected. For both of the above cases, not all packet sizes are affected. Packets less than 128 bytes in size are not affected, while packets above 128 bytes in size are affected at different packet size boundaries. For both of the above cases, messages similar to the following will be reported by the "-ES" FPCs and logged in the system messages file. Sep 26 14:36:47 routername fpc7 SRCHIP(0): 71024 Bad packets on p1 Sep 26 14:36:47 routername fpc7 SRCHIP(0): 71815 SONN errors on p1 Sep 26 14:36:47 routername fpc7 SRCHIP(1): 71056 Bad packets on p1 Sep 26 14:36:47 routername fpc7 SRCHIP(1): 71826 SONN errors on p1 [[PR794978](#): This issue has been resolved.]
- When dynamic-profile versioning is enabled, CoA requests are not processed if authd restarts. [[PR796416](#): This issue has been resolved.]
- In subscriber management environment, with GRES enabled, subscriber management infrastructure daemon (smid) maintains a directory /mfs/var/sdb which contains several logs/stats and databases that must be cleaned up and compacted every couple of hours. If smid process compacts databases during berkeley database replication daemon (bdbrepd) and has not finished the initial replication between Routing Engines, smid might get stuck in a loop cause it does not checkpoint or archive the log files anymore, so the /mfs directory eventually fills up. Then the router locks up with no telnet/console access, and no subscribers logging into the router. When issue happens, the following errors could be seen: /kernel: Process (1977,bdbrepd) has exceeded 85% of RLIMIT\_DATA: used 114692 KB Max 131072 KB /kernel: Process (1977,bdbrepd) attempted to exceed RLIMIT\_DATA: attempted 131076 KB Max 131072 KB /kernel: pid 1977 (bdbrepd), uid 0 inumber 636230 on /mfs: filesystem full The storage usage of bdbrepd process could be observed by following command (the 'SIZE' field means Total size of the process (text, data, and stack), in kilobytes): user@router> show system processes extensive | match bdbrepd PID USERNAME THR PRI NICE SIZE RES STATE TIME WCPU COMMAND 1568 root 1 96 0 55068K 37940K select 28:32 0.00% bdbrepd [[PR796430](#): This issue has been resolved.]
- This PR provides fix for a Packet Forwarding Engine micro-kernel memory leak, applicable only to certain types of firewall filter configurations on MX Series routers

with MPCs and DPCs. The concerned leak can occur with dynamic-profiles configuration with firewall filters (typical for subscriber services), or when protocols like Ethernet OAM are configured. The leak occurs only under a certain timing condition (so far observed for about 2-5% of interface delete operations), and only on interface delete operations. Leaked memory is of order of 1KB per incidence of the memory leak. Hence the issue will be of material impact in and only in deployment scenarios which involve many thousands of interface delete operations with either dynamic profiles configuration or Ethernet OAM configuration. [PR797790: This issue has been resolved.]

- It has been observed in the test lab that when there is a prefix and pointing to a multipath BGP nexthop (8 in number), and in turn each of this next-hops are pointing to multiple MPLS LSPs, the convergence number for 450k routes were in the order of 15 to 20 minutes. It is also observed that any change in nexthop, such as the interface flapping or neighbor flapping, had a significant impact on convergence time computing the convergence for this 450 k prefix routes \* 64 nexthop (8 paths for each of nexthop). [PR798771: This issue has been resolved.]
- In subscriber management environment with knob versioning enabled for dynamic-profile, during subscriber flapping or login/logout, the authd (authentication daemon) process may leak memory. [PR800028: This issue has been resolved.]
- On M Series and T Series platforms, in L3VPN scenario with l3vpn-composite-nexthop enabled, while PE receiving packet with DF bit set and packet size is larger than the core-facing interface's MTU, FPC/FEB may crash and generate a core file due to software defect of handling composite nexthop which needs packet fragment. [PR800155: This issue has been resolved.]
- When there are multiple recursive routes available for a prefix such as BGP route pointing to a indirect (8) and in turn pointing to a unilist (32) and when this is going over an AE of 8 links, we saw the software going through a large computation for each prefix and this makes it worse proportional to number of prefixes. At this instance, the system crashes because of the large computation. [PR800157: This issue has been resolved.]
- Dynamic VLAN might remain stuck in "terminating" state after line card reboot. [PR800533: This issue has been resolved.]
- [MPC] Non-QX cards do not transmit PPP hellos on wire. [PR801565: This issue has been resolved.]
- If MPC3 is equipped with 10x10GE MIC or 2x40GE MIC in the MIC slot 0 and 20x1GE MIC in the MIC slot 1, the links will not come up for MIC in MIC slot 0. [PR803613: This issue has been resolved.]
- In a BRAS environment with PPPoE subscribers and parameterized firewalls and policers, the FPC memory usage will go high after a CoA change for the services' parameters. [PR805922: This issue has been resolved.]
- Problem description: ----- The KSYNCD core followed by kernel live core is observed very rarely after the Routing Engine switchover. How to detect the problem? ----- This issue can be detected when ksyncd core is observed along with below mentioned log message in /var/log/messages. "Aug 27 01:28:03 indiranagar1 ksyncd[2506]: KSYNCD: resync error, issu\_state[0], type

Generic config subtype 8 : File exists" How to recover from the problem?

----- - Reboot the backup Routing Engine. [[PR810787](#): This issue has been resolved.]

- The ATM MIC with SFP (model number: MIC-3D-8OC3-2OC12-ATM) might show 0 pps values while executing **monitor interface at-x/x/x** even though there is input traffic on the interface's logical units. This not a service impacting issue. [[PR815632](#): This issue has been resolved.]
- In subscriber management scenario, during DHCP or PPPoE subscribers login, cosd process may leak memory when cosd is initializing cos related parameters for subscribers. It happens only if the subscriber has one cos service session. [[PR815777](#): This issue has been resolved.]

### **Infrastructure**

- On J Series devices, the top utility with "ores" options was not sorting the output based on resident memory size. [[PR507675](#): This issue has been resolved.]
- Under certain rare conditions Kernel "devfs" may become locked which may cause other processes that use /dev/filesystem to wait. Eventually some processes start spawning until reaching the maximum limit; as a consequence the kernel will crash. The following message logged by the kernel is an indication that the system is approaching the maximum number of active processes: /kernel: %KERN-2: nearing maxproc limit by uid 0, please see tuning(7) and login.conf(5). /kernel: %KERN-2: Process with Most Children- 1:init - Children - 365 [[PR678971](#): This issue has been resolved.]
- Ethernet driver for the internal Ethernet interface on the Routing Engine causes the kernel to crash and the Routing Engine reboot. This problem only could happen on the Routing Engine models that use "bcm" type of Ethernet interfaces for internal communication. To identify if the Routing Engine is using this type of interface, use the following command: user@router-re1> show interfaces terse Interface Admin Link Proto Local Remote [.....] bcm0 up up <----- bcm0.0 up up inet 10.0.0.1/8 10.0.0.5/8 128.0.0.1/2 128.0.0.5/2 inet6 fe80::201:ff:fe00:5/64 fec0::a:0:0:5/64 tnp 0x5 [.....] lsi up up mtun up up pimd up up pime up up tap up up [[PR734419](#): This issue has been resolved.]
- When the delete command is issued on an unnumbered Ethernet user route (static route with a qualified next hop ), the destination route created as a part of this user route does not get deleted. This results in duplicate ARP entries for the same address. [[PR752163](#): This issue has been resolved.]
- The Junos OS kernel might crash because of a timing issue in the ttymodem() internal I/O processing routine. The crash can be triggered by simple remote access (such as Telnet or SSH) to the device. [[PR755448](#): This issue has been resolved.]
- When Layer 3 VPN instances with localization were deleted/added multiple times, ip addresses were not created properly for interfaces. This issue is fixed now. [[PR769591](#): This issue has been resolved.]
- If a router is configured with a POSIX-compliant time zone string, it does not update the time zone correctly. Problem can be observed in system logs and CLI commands

when date/time zone is referenced. set groups re0 system time-zone EST5EDT,M10.3.0/2,M2.3.0/2 The router will incorrectly reference the previously configured time zone. After time zone configuration modification run "commit full". [PR785946: This issue has been resolved.]

- If we flap the interface used as next hop in the forwarding table for the IPv6 remote router loopback address used for IPv6 BGP sessions, the session flaps although there is another valid route over the other interface. [PR791881: This issue has been resolved.]
- Filter Based Forwarding is not working for IPv6 traffic [PR795730: This issue has been resolved.]
- In IPv6 scenario, when "ipv6-duplicate-addr-detection-transmits" is configured with a value of zero, IPv6 Neighbor Discovery might not function properly. [PR805837: This issue has been resolved.]
- Router might reboot while running 'show system core-dump core-file-info' command. This command uses /tmp and while uncompressing the core file, /tmp file system might be exhausted. /tmp in turn uses swap device only. MFS (Memory File System) and the rest of OS share the same swap space. Consuming more swap spaces might lead to out of memory and swap situation, which could eventually bring down the system. [PR808243: This issue has been resolved.]

### ***Interfaces and Chassis***

- Regarding the K2-RE (64-bit Routing Engine) when speed/link mode are statically configured on the router for the fxp0 interface, the driver for fxp0 accepts the configuration from DCD process, but does not propagate the setting to the hardware driver. Instead, the driver setting is forced to auto-negotiate. Thus, as the fxp0 interface is auto-negotiating, and the far end device is forced to 100/full, the auto-negotiation on fxp0 will detect the speed but not the duplex and defaults that duplex to half-duplex. [PR704740: This issue has been resolved.]
- After an MS-DPC generates a core file, it might result in I-chip "stream blocked detected". Traffic flow will be dropped and can only be restored by restarting the DPC. This is due to SG FPGA soft reset issue. [PR743262: This issue has been resolved.]
- There can be a mismatch between the ifIndex value on IF-MIB-ifName and the ifIndex value on SONET-APS-MIB-apsMapGroupName and apsMapEntry. [PR771877: This issue has been resolved.]
- Junos OS does not support unified in-service software upgrades (unified ISSU) for configurations that include interface sets. [PR-779377] [PR779377: This issue has been resolved.]
- When you issue the "show vrrp brief" command, a VRRP process (vrrpd) core file might be created. [PR782227: This issue has been resolved.]
- Load average values collected via SNMP do not show the correct values of the other Routing Engine. This can be verified by using the following commands: show snmp mib walk jnxOperatingEntry | match LoadAvg.9.1.0.0 show snmp mib walk jnxOperatingEntry | match LoadAvg.9.2.0.0 [PR782817: This issue has been resolved.]

- The prefer-status-control-active configuration knob at [edit interfaces aeX aggregated-ether-options mc-ae events iccp-peer-down] hierarchy requires configuration knob to be active at the [edit interfaces aeX aggregated-ether-options mc-ae status-control] hierarchy. When this is not present prefer-status-control-active has no impact and its presence in the configuration knob wrongly implies that the current node is preferred active. [[PR785930](#): This issue has been resolved.]
- In the environment of composite-next-hop with FMBB (fast make-before-break) function (e.g. VPLS, multicast, P2MP LSP), if system is configured with feature which contains interface having active/standby links (e.g. RLSQ/AMS/RMS), the Packet Forwarding Engine having standby interface might be incorrectly involved when building Packet Forwarding Engine flooding tree. This results in traffic blackhole. [[PR786007](#): This issue has been resolved.]
- 'monitor ethernet delay-measurement' command does not timeout when CFM adjacency is down and/or all DMM frames are sent. As a result, ethdm binary does not close normally leading to an increase in resource consumption. [[PR787985](#): This issue has been resolved.]
- FPC might crash due to page fault [[PR791195](#): This issue has been resolved.]
- T640 frame relay interface status is shown as up/up with mismatched lmi-type. [[PR791501](#): This issue has been resolved.]
- Configuring fxp0 with Speed 10m and Full-Duplex generates log message "fxp0: Full duplex link mode is not supported with speed 10M, Hence Speed will default to 100M" [[PR791777](#): This issue has been resolved.]
- "show chassis hardware" output for some optics is not correct sometimes. [[PR792704](#): This issue has been resolved.]
- On MX Series routers, a change to the 'oam lfm pdu holdtime' on an interface is not updated correctly. This results in an incorrect LFM state, which should be reported as Adjacency Lost. As a workaround, issue the clear oam ethernet link-fault-management state command from the CLI to correctly update the 'pdu holdtimer.' [[PR792763](#): This issue has been resolved.]
- Issue is seen only when the following steps are followed: 1. Enable IRB MAC Sync feature. 2. Deactivate BD/MAC Sync/Service ID on the higher MAC node. 3. Activate virtual switch that configures MCAE under the virtual switch. The result: IRB MAC Sync happens even though the feature is not enabled. [[PR793889](#): This issue has been resolved.]
- When you configure the untagged GE interface or untagged aggregate Ethernet with link member on IQ PIC with per-unit-scheduler, this might cause the failure of interface statistics on this interface. As a result the following error will be reported in the log message and on "show interface extensive" outputs. [[PR794975](#): This issue has been resolved.]
- Continuously walking SNMP MIBs while PPP subscribers log in and out of the router causes the PPPoE daemon to go unresponsive due to synchronous retrieval of SNMP MIBs take larger time and blocking signals to be serviced by PPPoE daemon. Subscribers can no longer log into the router. The daemon will remain unresponsive as follows until it is restarted. user@router> show pppoe statistics error: the pppoe subsystem is not

responding to management requests The SNMP MIBs being walked are as follows:

1.3.6.1.4.1.2636.3.68.1.1 (jnxPPPOObjects) 1.3.6.1.4.1.2636.3.67.1 (jnxPPPoEMIB)  
 1.3.6.1.4.1.2636.3.64.1.1 (jnxSubscriberObjects) 1.3.6.1.4.1.2636.3.12.1.1.1 (jnxIpv4AddrEntry)  
 1.3.6.1.4.1.2636.3.51.1.1 (jnxUserAAAOObjects) [[PR795556](#): This issue has been resolved.]

- When upgrading to 11.4R4, links that are using tunable DWDM XFP are not working anymore and are reporting a different wavelength than the configured one. [[PR796330](#): This issue has been resolved.]
- jpppoed memory utilization spikes after GRES or jpppoed restart event. [[PR800650](#): This issue has been resolved.]
- An operation in order of PIC offline then deactivate ci member interface will have the deactivated member interface join to the ci upon PIC online. And it will cause unexpected forwarding issue. [[PR803817](#): This issue has been resolved.]
- When a hold-time is configured for xe interfaces, before the timer is started, the XFP alarm is checked before declaring the link status. As a result, it is possible for the link to remain down if there are XFP alarms/warnings as reported by the "show interfaces diagnostics optics" command. [[PR804315](#): This issue has been resolved.]
- In 11.4R4, 12.1R3, 12.2R1, and later, the option of 'routing-engine' under "> request system snapshot" was mistakenly removed. [[PR809321](#): This issue has been resolved.]
- When the Flexible PIC Concentrator (FPC) restarted after performing a master Routing Engine switchover, the aggregate interface flag was set to down. Any traffic that entered this FPC and traversed the equal-cost multipath (ECMP) to the aggregate interface was not using the aggregate member interfaces. For vpls/bridge traffic the effect is on the ingress side processing mac address learning. Since the aggregate interface is down, mac address processing is terminated and re-learned again once packets are received. This triggers high CPU utilization of the FPC linecard hosting those member links. [[PR809383](#): This issue has been resolved.]
- Interface damping is not working properly on MPC Type 2 3D [[PR810159](#): This issue has been resolved.]
- VC powerdown of VCMm leaves lp demux interfaces in a "hardware-down" state [[PR813902](#): This issue has been resolved.]

### **Layer 2 Features**

- If a VPLS interface is moved to a VPLS aggregate bundle or changed to any other interface family different than vpls, on the next GRES/nonstop active routing (NSR) Mastership switch, this interface will have the CCC-Down flag set and will not process any traffic. The interface needs to be deactivated and activated via configuration change to continue forwarding traffic. [[PR788631](#): This issue has been resolved.]
- After a physical interface that is configured for CoS on a logical tunnel interface flaps, the MAC address of the peering logical interface goes missing from the kernel. To resolve this issue, deactivate/activate the peering logical interface after the flap. [[PR790559](#): This issue has been resolved.]



- "monitor traffic interface irb" fails to capture outgoing packets. [[PR802605](#): This issue has been resolved.]
- Routing Engine CPU utilization may increase faster than expected when LDP neighbors are configured under a mesh group of a local BGP-VPLS routing instance when the LDP neighbors themselves are not provisioned for VPLS service on the remote side. [[PR808333](#): This issue has been resolved.]

### ***Layer 2 Ethernet Services***

- LACP status disagreement after the Routing Engine switchover. [[PR751745](#): This issue has been resolved.]
- DHCP server in a vrf responds with incorrect server identification address option 54 in DHCPOFFER. [[PR776222](#): This issue has been resolved.]
- The AFTR information from RADIUS is advertised by MX Series router to the client via DHCPv6. [[PR779679](#): This issue has been resolved.]
- Old dhcp session still can be renewed by a client after the client moves to another VLAN. [[PR784951](#): This issue has been resolved.]
- In an MX Series Virtual Chassis environment, a traffic outage longer than 2 minutes occurs when a member of the VC-M LAG fails while LACP is in active mode with link protection on the VC-M router. The outage occurs while the LACP process restarts on the new VC-M router. To avoid this situation, make sure that LACP is running in active link protection mode on the device in front of the VC-M router. This device cannot be an EX Series Ethernet Switch, because the switch does not support LACP in link protection mode. [[PR784965](#): This issue has been resolved.]
- In 11.4R2, 12.1R1, and 12.2R1 and then subsequent builds on those releases, there might be a false alarm of a hardware problem from a DPC. For example: "fpc0 EZ: %PFE-3: ezchip\_periodic\_check\_free\_rfd\_buffer[4245] XETH(0/3) : Rx RFD buffers exhausted" This can be ignored, unless traffic impact is seen. [[PR796824](#): This issue has been resolved.]
- DHCP relay does not forward ACK to client from the backup DHCP server after primary DHCP server failure. [[PR799090](#): This issue has been resolved.]

### ***MPLS***

- Configuration changes to some attributes of the standby secondary path of an MPLS label-switched path (LSP) might cause the LSP to flap, which might result in packet loss. [[PR394184](#): This issue has been resolved.]
- In GRES (graceful Routing Engine switchover) mode and due to a quick status change of MPLS CCC nexthop, a mismatch of index value between master and backup Routing Engines might happen, causing the backup Routing Engine to panic, generate a core file, and trigger a live core dump from the master Routing Engine. [[PR755473](#): This issue has been resolved.]
- The issue appears to happen when MVPN tries to add a vt-interface for an egress tunnel. RSVP would try to find the flood nexthop for the route installed for the label for the branch LSPs in the P2MP LSP to add the vt-interface. When RSVP does not



find the flood nexthop for the label route for a branch LSP, it triggers an assertion failure. [[PR770538](#): This issue has been resolved.]

- When you commit a configuration change that simultaneously disables RSVP and a point-to-point interface (so, t1, atm), an rpd core file might be generated. To solve this issue, do not commit a configuration change that simultaneously disables RSVP and a point-to-point interface. Rather disable RSVP and point-to-point interfaces in separate configuration commits. [[PR782174](#): This issue has been resolved.]
- With the fix of this PR, at the end of the adjust-interval operation, the max\_average counter does not reset to zero and maintains the old value until it receives a new sample. When the new sample is received, the max\_average counter is updated to the new sample value. This is a just display counter fix and there is no operational impact. The auto-bandwidth functionality works as it is. [[PR799155](#): This issue has been resolved.]
- Point-to-multipoint inclusive tunnels over MVPN might not come up because the RSVP state for the "vt" logical interface appears as down. [[PR802344](#): This issue has been resolved.]
- Some implementations of cSPF allow zero-bw LSPs (LSPs, which are requesting bandwidth of 0bps) to be calculated via links which have 0 bps of available bandwidth. In some cases implementation of RSVP in Junos OS allows AvailBW on link to become negative. This might happen, for example, in case of failure of one of links in ae bundle. As it's impossible to include negative values in OpaqLSA, in this case Junos OS announces 0 bps of AvailableBW on such links. Zero-bw LSPs will not be established via such link, because transit node with negative AvailBW fails check for bandwidth availability on egress interface. HeadEnd will constantly try to signal LSP and every time it will receive PathErr: BW Unavailable from node which has link with negative AvailBW. cSPF calculation will not resolve this, because according to TED on HEnd we have 0 bps of AvailBW on this link, not negative. [[PR802995](#): This issue has been resolved.]
- With RSVP disabled, when an SNMP get/get-next is received for RSVP MIB, a Path State Block (PSB) search request is queued. This enqueue operation returns nothing but the memory allocated for the search request is not freed and this results in a memory leak of routing protocol process (rpd). The memory leak could be observed by the following commands: user@router> show task memory detail | match "rsvp psb lookup req" ----- Allocator Memory Report -----  

Name	Size	Alloc	DTP	Alloc	MaxAlloc	MaxAlloc	Size	Blocks	Bytes	Blocks	Bytes
RSVP PSB lookup req	176	180	T	110	19800	110	19800				

 user@router> show system processes extensive | match rpd  
 PID USERNAME THR PRI NICE SIZE RES STATE TIME  
 WCPU COMMAND  
 1311 root 1 4 0 1529M 1479M kqread 75:25 0.44% rpd  
 When the memory usage of rpd process increases to around 85% of system limit, the following logs could be seen: re0: /kernel: %KER-5:Process (1859,rpd) has exceeded 85% of RLIMIT\_DATA: used 1835088 KB Max 2097152 KB [[PR811951](#): This issue has been resolved.]
- RSVP interface state will show as DOWN on unnumbered p2p interfaces. This will cause MPLS LSPs over those interfaces to stay down. [[PR814071](#): This issue has been resolved.]

**Network Management and Monitoring**

- After a Routing Engine switchover, LACP and MIB process (mib2d) core files might be created. [[PR790966](#): This issue has been resolved.]
- The security name specified under 'target parameters' should be used while sending v3-specific v1/v2 traps. [[PR813430](#): This issue has been resolved.]

**Platform and Infrastructure**

- You might see zombie processes increment while doing commit each time [[PR692382](#): This issue has been resolved.]
- Under very special race conditions the MPC CPU might stop processing and will be reset due to Level3/Level 2 watchdog expiration timer. Potential exposure is high load of traffic sent to the Host. The following syslog message will be reported in the syslog once MPC reboots. "fpc[x] MPC: Reset reason (Oxc): Level3 watchdog, Level2 watchdog" [[PR717899](#): This issue has been resolved.]
- PTSP and AACL services do not work with AMS interfaces. [[PR727588](#): This issue has been resolved.]
- In scenario where telnet session is disconnected ungracefully while accessing "load merge terminal" prompt problem can be exhibited with other CLI users unable to access configuration mode. [[PR745280](#): This issue has been resolved.]
- Memory exhaustion on the Packet Forwarding Engine ukern heap causing FPC core. [[PR777609](#): This issue has been resolved.]
- Need to add 'start-time', 'stop-time' and 'timezone' attributes in TACACS+ accounting packets. Solution: ----- Added a configuration knob to allow enabling these attributes to be framed into the accounting packets. Default behavior is to not include these attributes in the packet to maintain backward compatibility. The new knob is [ system tacplus-options timestamp-and-timezone ]. You will have to enable this knob and then check for the 'start-time', 'stop-time' and 'timezone' attributes in the accounting packets. [[PR780484](#): This issue has been resolved.]
- When MX Series routers with MPCs/MICs inline NAT translates a UDP pkt with UDP checksum 0x0000, it rewrites checksum to nonzero value. [[PR782927](#): This issue has been resolved.]
- Flow records obtained by using "inline-jflow" might contain incorrect AS value. Issue has been fixed in 11.4R6, 12.1R5, and 12.2R2 onwards. [[PR788879](#): This issue has been resolved.]
- When reconfiguring an interface from a native VLAN to another tagged VLAN, the logical interface mapping on the Packet Forwarding Engine might get corrupted. In case traffic is being received on this interface, it can lead to LU congestion and wedge. [[PR792633](#): This issue has been resolved.]
- Committing a Q-in-Q configuration results in an FPC crash in these conditions: 1. Core facing interface is configured this way: flexible-vlan-tagging; encapsulation flexible-ethernet-services; unit xxx { encapsulation vlan-bridge; vlan-tags outer xxx inner-range 1-4094; } 2. Core-facing interface is an aggregate interface. 3. Core-facing interface is on MPC card. [[PR793429](#): This issue has been resolved.]

- On MX Series routers with MICs/MPCs (in Releases 11.4R4+, 12.1R3+, 12.2R1+), when the first large-sized (>1500) transit packet hits a resolve route on the Packet Forwarding Engine, it might cause memory to leak on the Packet Forwarding Engine micro kernel. [[PR802051](#): This issue has been resolved.]
- In Junos OS Release 11.2 Layer 3 services over MC-LAG are supported through IRB only and thus family inet is not supported directly on MC-LAG interface. However, appropriate commit check is missing in 11.2R3 and later maintenance releases of 11.2 [[PR802938](#): This issue has been resolved.]
- With inline sampling, when there are multiple flow servers being configured or multiple equal cost paths exist for a single collector, the flow record packet might trigger the following trap message from the Packet Forwarding Engine which causes a drop for the flow record packet. PPE Sync XTXN Err Trap: Count 1659, PC 45f, 0x045f: balanced\_multi\_nh\_use\_cp\_index [[PR805061](#): This issue has been resolved.]
- IPv6 traffic on MX Series routers with MPCs/MICs hardware might cause IPv4 SCU/DCU counters to increment. [[PR805257](#): This issue has been resolved.]
- AIS scripts error "error: xsl:import : unable to load" [[PR815978](#): This issue has been resolved.]

### ***Routing Protocols***

- Dynamically signaled routes might flap when a **commit full** is performed. [[PR672838](#): This issue has been resolved.]
- RPD might core after deleting or renaming a non-forwarding instance. Specifically, issue occurs when: 1. the interface is configured in a non-forwarding instances (that is, routing instances xxx with no instance-type). and any one of the following: 2a. igmp is configured with all interfaces (e.g. protocols igmp interface all) 2b. igmp is configured on the specific interface (e.g. protocols igmp interface ge-0/0/0.1) 2c. mld is configured with all interfaces (e.g. protocols mld interface all) 2d. mld is configured on the specific interface (e.g. protocols mld interface ge-0/0/0.1) 2e. pim is configured in the master instance with all interfaces (e.g. protocols pim interface all) 2f. pim is configured in the master instance on the specific interface (e.g. protocols pim interface ge-0/0/0.1) and any one of the following actions are subsequently committed: 3a. delete the non-forwarding instance (e.g. delete routing-instances xxx) 3b. rename the non-forwarding instance (e.g. replace pattern xxx with yyy) [[PR704699](#): This issue has been resolved.]
- If there is more than one peer in a BGP group, and have nonstop active routing (NSR) configured, in a rare condition, BGP session init job fails on the last established peer in the group, while at the same time, there is another peer in establish-ack-wait state. Then after that, the peer in establish-ack-wait state becomes established after getting an establish-ack from the remote peer which will cause the routing protocol process (rpd) to crash and generate a core file. [[PR736198](#): This issue has been resolved.]
- RPD can crash soon after OSPF switches from primary path to secondary path when LFA (loop free alternates) is enabled, along with LDP-SYNC: /kernel: BAD\_PAGE\_FAULT: pid 1472 (rpd), uid 0: pc 0x86ff81c got a read fault at 0x15, x86 fault flags = 0x4 The corruption happens because of race condition, when OSPF doesn't completely free a memory location which is later reused by LDP. [[PR737141](#): This issue has been resolved.]

- If a routing instance is configured to add static routes to its instance-specific routing table using both the routing-options static route stanza and the routing-options rib <instance specific table name> static route stanza and a configuration event changes something else in the routing-options rib <instance specific table name> stanza such as modifying the maximum-paths value, the static routes in the instance table specific section can be deleted. A commit full can be used to recover or only use one of the 2 mechanisms for defining the static routes for that instance. [[PR755558](#): This issue has been resolved.]
- The rpd generates a core file at "rip\_dc\_retrans\_callback\_p2mp" while unconfiguring p2mp configuration. [[PR769487](#): This issue has been resolved.]
- In PPPoE subscriber management environment, without access-internal configured under [dynamic-profiles routing-options] hierarchy, after subscribers flapping, Access Internal route of PPPoE subscribers might lose. This condition causes traffic loss for the subscribers that have no access-internal routes installed. And if the access route is installed, it will be hphysical interface as there is no usable next-hop. [[PR772882](#): This issue has been resolved.]
- The current implementation of the random number generator grand() is not thread safe, as it utilizes a number of global variables and arrays. When bgp precision timers are enabled, this can cause crashes, e.g., when we jitter timers. The changes herein give each thread its own random number generator. [[PR777802](#): This issue has been resolved.]
- The rpd crashes when the customer specific BGP configuration is unconfigured. [[PR782816](#): This issue has been resolved.]
- On Junos OS platforms with 64-bit kernel (Routing Engines), a bug in the routing protocol process (rpd) memory allocator might cause corruption in the rpd memory which might finally lead to an RPD crash. [[PR792238](#): This issue has been resolved.]
- During extended stress testing of the PIM protocol, a malformed PIM Hello message triggered an RPD crash. While the crash was caused by a malformed PIM message, simply replaying the crafted packet alone does not lead to the crash. This issue affects both IPv4 PIM and IPv6 PIM. Refer to PSN-2012-10-732 for additional information. [[PR792334](#): This issue has been resolved.]
- An MX Series router which has only some bridge-domains configured for igmp-snooping might discard traffic in bridge-domains without igmp-snooping enabled. [[PR795781](#): This issue has been resolved.]
- Setting OSPF overload via the configuration sets both the metric field in router LSAs as well as te-metric field in opaque LSAs to 65535 or  $2^{16}-1$ . Since te-metric is a 32-bit field, it should be set to  $2^{32}-1$ . [[PR797293](#): This issue has been resolved.]
- When using precision-timers, it is possible that BGP does not send all its advertisements from its buffer, until the BGP session needs to send another non-keepalive message. [[PR801037](#): This issue has been resolved.]
- The RPD generates a core file after making changes to policy-statement related to VPN [[PR807357](#): This issue has been resolved.]

- The below log messages are generated when a commit is performed.  
"task\_set\_option\_internal: task ICMP socket 103 option GroupAdd(23) interface ae12.0: Address already in use." The error leading to this log is handled properly. Since these logs can be misleading the logs have been suppressed with the fix. [[PR809472](#): This issue has been resolved.]
- Unicast RPF check not working properly after bouncing the RPF interface, when there are two BGP routes for same destination via two different next-hop RPF interfaces. [[PR814303](#): This issue has been resolved.]
- RPD in backup cored@rt\_nexthops\_free multiple times [[PR816754](#): This issue has been resolved.]

### Services Applications

- When a TX router is configured with a manual OSPF ipsec-sa for authentication, something like the following cosmetic messages will be logged: Feb 16 16:27:40 flame-sfc-re1 lcc0-master kmd[17194]: KMD\_RTsock\_ERROR: Error adding inbound SA OSPF3\_AH\_SHA1\_96 spi=1024 proto=AH to kernel: No such file or directory Feb 16 16:27:40 flame-sfc-re1 lcc0-master kmd[17194]: KMD\_RTsock\_ERROR: Error adding outbound SA OSPF3\_AH\_SHA1\_96 spi=1024 proto=AH to kernel: No such file or directory If there's a service PIC, there will be these additional cosmetic log entries: Feb 16 16:27:46 flame-sfc-re1 lcc1-master kmd[16853]: KMD\_INTERNAL\_ERROR: Failed to connect PIC, ERR: Failed to connect PIC, ERR: F Feb 16 16:27:46 flame-sfc-re1 lcc1-master kmd[16853]: KMD\_INTERNAL\_ERROR: Unable to connect PIC sp-8/3/0; Feb 16 16:27:46 flame-sfc-re1 lcc1-master kmd[16853]: KMD\_INTERNAL\_ERROR: Couldn't request PIC: sp-8/3/0 to send sa state [[PR738736](#): This issue has been resolved.]
- This crash was observed during a mixed traffic test for 7 hours on below traffic profile. The traffic profile used HTTP 0.8m HTTPS 0.15m FTP 0.1m RTSP 0.08m UDP 8.87m (IMIX traffic) [[PR769322](#): This issue has been resolved.]
- Memory leak on FPC/FEB when service next-hops are deleted on LNS for l2tp sessions on non-MX Series platforms. For each l2tp session teardown, there are 240 bytes of memory leaks. After some time this leak will lead to FPC/FEB memory exhaustion leading to unpredictable FPC/FEB reset. [[PR770903](#): This issue has been resolved.]
- RTSP streaming is not working in a laptop when moving or fast forwarding the video. [[PR786085](#): This issue has been resolved.]
- In L2TP setup with MX Series router acting as LAC, if the value used passed from RADIUS in VSA "Tunnel-Client-Endpoint" does not exist on the router, Junos OS will send SCCRQ message to LNS with random source addresses. [[PR788081](#): This issue has been resolved.]
- When an MX Series router configured as an LNS sends an Access-Request message to RADIUS for an LNS subscriber, the LNS now includes the Called-Station-ID-Attribute when it receives AVP 21 in the ICRQ message from the LAC. [[PR790035](#): This issue has been resolved.]
- MX LNS does not support CLI command services l2tp session user filter option. [[PR792239](#): This issue has been resolved.]

- The **clear services l2tp session user <user-name>** command accepts any arbitrary alphanumeric characters in place of a user name and the CLI command will drop all l2tp subscribers. [[PR792631](#): This issue has been resolved.]
- IDP process might go down temporarily if multiple IDP detectors are installed. [[PR794335](#): This issue has been resolved.]
- [Deterministic NAT]ports allocation overlapped [[PR797457](#): This issue has been resolved.]
- KMD process running high with key chain configuration. Need to modify the necessary changes so that KMD does not run wherever it is not required. [[PR798030](#): This issue has been resolved.]
- The MX Series CLI allows you to configure "\*" for client name in the l2tp access profile leading to failure of establishing the l2tp connection. [[PR799232](#): This issue has been resolved.]
- deNAT:wrong mapping between nat-port-block and internal-host [[PR799947](#): This issue has been resolved.]
- In scenario where the MX Series router is acting as LAC and radius server is returning tunnel-server-endpoint attribute but NOT returning tunnel-client-endpoint, memory leak in jl2tpd process can occur. Additionally same memory leak can occur if unnumbered loopback attribute is returned from radius for tunneled subscribers. [[PR800107](#): This issue has been resolved.]
- In a carrier-grade NAT (CGN) configuration that includes the **address-pooling paired** statement (APP feature) the service PIC might reset unexpectedly. This behavior occurs only in certain corner case scenarios (with low probability to be seen in production) in which the service PIC software times out in an APP mapping but a new flow is created within the same mapping. There is no workaround. [[PR800241](#): This issue has been resolved.]
- MXVC-L2TP Core /src/bsd/lib/libthr/thread/thr\_kern.c:91 [[PR802044](#): This issue has been resolved.]
- NAPT:port range starts from 512. It should start from 1024. [[PR804598](#): This issue has been resolved.]
- When adding/deleting/changing the address ranges configured under NAT pool, the Service PIC might run into an inconsistent state and restart. [[PR810994](#): This issue has been resolved.]
- Refer to the AT, for the list of new commit constraint checks imposed. [[PR815053](#): This issue has been resolved.]
- Internally, each term is treated as a rule. Multiple copies of NAT pool are created, one for each rule (or term). When address-allocation round-robin is configured, the NAT ip is read from separate copy of the NAT pool; thus we see only 10 NAT IPs allocated

for 20 different hosts matching two different terms. After the fix , all terms now can share the same NAT pool. [[PR815147](#): This issue has been resolved.]

- When NAT pool mapping for a private address is being deleted, and at the same time there is a flow being created for the same private address, then service PIC might crash in rare cases. [[PR821037](#): This issue has been resolved.]

### ***Subscriber Access Management***

- A Change-of-Authorization request is being NAK-ed on MX80 when a PPP subscriber is terminated in a non-default routing-instance. It works as expected if the subscriber is terminated in the default routing-instance. [[PR704560](#): This issue has been resolved.]
- After a large number of concurrent PPP session logouts and GRES operation, some sessions might not complete logout (services activated from SRC). Sessions eventually will time out and clear. [[PR742900](#): This issue has been resolved.]
- Authd attempts to remove VLAN when subscribers are idle but connected [[PR789009](#): This issue has been resolved.]
- The captive portal content delivery service applied on PPPoE subscriber to rewrite IPDA is not working. The subscriber traffic is altered but is dropped on MS-DPC. [[PR789368](#): This issue has been resolved.]
- In subscriber management scenario, if the subscribers (such as PPPoE, DHCP, IPOE) log in through authenticated dynamic VLAN, after they log out, the authenticated dynamic VLAN will be removed, but the corresponding libstats iflstats entry will be left which should not happen. Finally the memory leak will cause filesystem /mfs to be full and prevent subscribers from logging in. If this issue happens, the following logs could be seen: /kernel: ifl(pp0.1073859148): ifl\_config not found !!! smid: /mfs capacity 83% smid: /mfs capacity 85% authd[13198]: ===== Idle Timeout Exceeded Rid the subscriber ===== last message repeated 2409 times smid: /mfs capacity 108% /kernel: pid 13197 (jdhcpd), uid 0 inumber 24280 on /mfs: filesystem full jdhcpd: DH\_SVC\_LOGIN\_FAILURE: DHCP pre-authentication failure for DHCPv4 client SDB session 8659554 on incoming interface demux0.1073841747 [[PR796299](#): This issue has been resolved.]
- MX-VC:Authd keeps retrying the attempts to fetch final stats for already removed logical interface. [[PR806104](#): This issue has been resolved.]
- MX-VC:Authd sends duplicate requests to enable interim accounting in PFED for idle timeout configured on VLAN subscriber [[PR806112](#): This issue has been resolved.]
- MX-VC:Authd core:../../../../authd\_aaa\_dyn\_req.cc:780 [[PR806128](#): This issue has been resolved.]
- Stale addresses in local address pool [[PR815331](#): This issue has been resolved.]



### VPNs

- An optimization has been implemented with BGP-MVPN nexthop infrastructure which will improve scalability in some multi-dimensional scaling scenarios with aggregate interfaces. [[PR690690](#): This issue has been resolved.]
- When you disable protocols in a Layer 2 circuit with egress protection, rpd generates a core file if no routes are found in context routing table. [[PR735789](#): This issue has been resolved.]
- UMH selection should select the highest IP address as the Upstream PE. However, in the code the highest IP address is selected by comparing lowest order byte of the IP address first. In this case between IP address 10.233.38.34 and IP address 10.233.32.46 - 10.233.32.46 gets chosen as upstream PE because its lowest order byte (46) is more than the lowest order byte of 10.233.38.34. This is because code does not account for the endian-ness of the machine it is run on. Fix - convert the IP address to network order before comparing. [[PR754114](#): This issue has been resolved.]
- The issue is seen when multicast traffic is stopped and PIM states are allowed to clear. It was seen that some of (S, G) states on the PE did not clear. As a result data streams for the affected groups do not reach the intended receivers. However it was seen that the fix for this issue caused another issue to appear as documented in PR# 823884 [In NG-MVPN RPT-SPT mode, failure and then recovery of the primary path to the RP, causes the PE routers to forward traffic on (\*, G) even though the (S, G) join is pruned] [[PR779786](#): This issue has been resolved.]
- In L2VPN scenario with at least one L2VPN connection in up state, during SNMP walk on jnxVpnPwAssociatedInterface.bgpL2Vpn (OID: .1.3.6.1.4.1.2636.3.26.1.4.1.6.3.9), an infinite loop occurs due to software defect and causes routing protocol process (rpd) to be stuck at 97% indefinitely until the rpd process is restarted. While in this state all protocol adjacency will expire and the router will stop forwarding traffic. [[PR782654](#): This issue has been resolved.]
- When the Source PE is rebooted, S,G,RPT state is not cleaned up. The change is to clean up this S,G RPT state if we learn the source address is remote. [[PR784627](#): This issue has been resolved.]
- In Rosen MVPN scenario, after performing the Routing Engine switchover, some of the Rosen VRFs don't have a tunnel interface (mt-) assigned for incoming, and therefore will not join the transport group sourced with the remote PE loopback address. [[PR791333](#): This issue has been resolved.]
- The rpd incorrectly sets the PWE3 Control Word flag for local switching circuits. PWE3 Control Word is needed for Layer 2 VPN OAM packets to get pushed to the Routing Engine. On MX Series routers with MPCs/MICs, the traffic payload is examined and if it matches the first nibble being 0001, it sends the traffic to the Routing Engine for further processing. Once Layer 2 VPN traffic is set to IPv4, it would test against the IPV4 ID field. [[PR793751](#): This issue has been resolved.]
- This issue was experienced in Junos OS Release 11.4R4 code with NG-MVPN RPT-SPT mode. When a link failure causes the route to the source and RP via the backup path, the PE in the backup path fails to forward multicast traffic to the receiver. This is



documented in PR 794222 and upgrading to Junos OS Release 11.4R4-S2 fixes the above mentioned issue. [[PR794222](#): This issue has been resolved.]

- This change would allow customers to use the less restrictive CLI knob 'vrf-advertise-selective', which now accepts a null list. If no family is configured under vrf-advertise-selective, then no MVPN routes are advertised to the neighbor. [[PR795108](#): This issue has been resolved.]
- When the egress PE receives type 4 route before discovering the ingress PE, it queues all the type 4 routes for later processing. Due to a defect in the corresponding code, the replaying of type-4 was not happening always, thereby causing the ingress PE to not send the corresponding stream to the egress PE. [[PR801437](#): This issue has been resolved.]

#### **Release 11.4**

The following issues have been resolved since Junos OS Release 11.4R4. The identifier following the description is the tracking number in our bug database.

- [Forwarding and Sampling](#)
- [General Routing](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Layer 2 Ethernet Services](#)
- [MPLS](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)
- [Services Applications](#)
- [User Interface and Configuration](#)
- [VPNs](#)

#### **Forwarding and Sampling**

- When the configuration archiving FTP process stalls during file transfer, it can result in the PFED process stalling as well. Once the master PFED process is restarted, it results in the inability to commit certain new configuration changes. Ensuring that the configuration archiving and FTP server are correctly configured and working will avoid this problem. [[PR528653](#): This issue has been resolved.]
- Sampled memory increases when interfaces bounce and BGP is running. [[PR594509](#): This issue has been resolved.]
- On ADPC cards Output Layer-2 policer drops the packets when configured on an interface with vpls encapsulation. [[PR749141](#): This issue has been resolved.]
- Any change to the last member of a service-filter chain can lead to the loss of Layer 3 connectivity over the interface. [[PR750957](#): This issue has been resolved.]

- This issue can occur for daemons that connect to pfed for statistics information. If the daemon starts before pfed, or has problems making a connection to pfed, then that daemon could experience a crash. This can also occur using CLI commands such as **show interface statistics** that invoke ifinfo. This problem was introduced by the fix for [PR743135](#) and only exists in the specific releases that were fixed by that PR. There is no known workaround to this problem except to use a release of Junos OS with the fix. [[PR770766](#): This issue has been resolved.]
- This PR eliminates an erroneous error message that would appear in syslog while pfed was checking the syntax of a configuration containing firewall filters that reference accounting counters. This problem only affected the syntax check prior to the commit. The actual configuration on the router was correctly committed. [[PR772463](#): This issue has been resolved.]

### **General Routing**

- There is an issue with the interworking of 'chassis route-memory-enhanced' knob and 'protocols rsvp interface x/y/z.l link-protection' or 'protocols mpls label-switched-path <lsp name> node-link-protection'. This causes failure of the installation of routes that are destined to go to segment 1 (due to the route-memory-enhanced) configuration. [[PR695336](#): This issue has been resolved.]
  - Once a child interface of an aggregate bundle is in down state (for example: CCC-Down of the logical member link interfaces), the next-hop of the control channel is not correctly programmed. LACP packets received are not dropped but processed and point to invalid NH entries, which might yield to such errors as follows or a combination of all:
    - fpc5 LUCHIP(0) IDMEM[0x000433ba] Read Uninitialized Memory Error
    - fpc5 LUCHIP(0) PLCT INT\_STAT 0x00000001 Illegal PL Uninitialized EDMEM Read 0x6db6db6d6db6db6d @ 0x1cf30001 XTXN 0xa8cd87 BULK 0x005c0094 FN 0 sync PPE 14 CNTX 1
    - fpc5 LUCHIP(0) RMC 2 Uninitialized EDMEM[0x1001c0] Read (0x6db6db6d6db6db6d)
    - fpc5 LUCHIP(0) PPE\_6 Errors sync xtxn error thread timeout error
    - fpc5 PPE Sync DMEM WP Trap: Count 103, PC 620b, 0x620b: nat46\_loop 0x620b: nat44\_loop
    - fpc5 PPE Sync XTXN Err Trap: Count 980053, PC 2f9, 0x02f9: nh\_ret\_simple\_last
    - fpc5 PPE Thread Timeout Trap: Count 2840, PC 4c6, 0x04c6: set\_iif\_inc\_ifl\_cnt fpc5 PPE PPE Stack Err Trap: Count 20347, PC 310, 0x0310: add\_default\_layer1\_overhead
    - fpc5 PPE PPE HW Fault Trap: Count 529, PC 373, 0x0373: inner\_rewrite
- There is no operational impact other than the filling up of error messages in the system log. [[PR703245](#): This issue has been resolved.]
- If you have DCU statistics configured in conjunction with the copy-plp class of service knob and an output firewall filter, you might encounter a situation where DCU stats are no longer working. Check first that both ingress and egress ports for the flows you

are counting are not on different Packet Forwarding Engines the same fpc. If this is the case, then removing the copy-plp knob will restart the DCU statistics collection.

[[PR707834](#): This issue has been resolved.]

- On Systems platforms M320 E3FPC/M120/M7i(10i) CFEB-E with l2vpn or l2circuit, using a control-word and the mpls payload is corrupted in a certain way: the interface might stop forwarding traffic. To recover from this condition an FPC reboot is needed. Only Junos OS Release 10.0 or later is affected with non-cookie based PICs. MX Series platforms with DPC are not affected. [[PR720523](#): This issue has been resolved.]
- When receiving large bytes of PIM Join/Prune refreshes at a very rapid rate, it might exhaust ukernel buffer memory on the Packet Forwarding Engine, and the PIM Join/Prune packets will be lost. [[PR720966](#): This issue has been resolved.]
- The routing protocol process might generate core files when a community named in the policy options configuration is changed and that community-name is used in a **show route** command before the changes effectively take place. [[PR740427](#): This issue has been resolved.]
- When "delete", or "deactivate" "interface <unit>:family inet:accounting" configuration; FPCs that have the configuration removed might be reset unexpectedly. [[PR743442](#): This issue has been resolved.]
- In Junos OS Release 11.2R3 or earlier, if IPv6 traffic needs to trigger an ICMPv6 MTU exceeded message to the source and the source is resolved via next-table next-hop, it might leak packet memory on the FPC. [[PR745988](#): This issue has been resolved.]
- Jtree memory leak occurs on I-chip based platforms when 'route-memory-enhanced' or 'memory-enhanced' is enabled, after there is route flapping which is using indirect next hop. [[PR751567](#): This issue has been resolved.]
- When there are at least three routers to a specific destination (for example, two destination routes and one clone route), deleting and re-adding one of the logical interfaces (for example, board replacement) might trigger a kernel crash due to a timing issue with route deletion. This is triggered in the specific topologies such as an OSPF3 next hop, which is connected to a different vendor device:

```
lab@shark-re0> show route forwarding-table destination
fe80::21f:9eff:fea9:c140
Routing table: default.inet6
Internet6:
Destination      Type RtRef Next hop    Type Index NhRef Netif
fe80::21f:9eff:fea9:c140/128      dest      0 0:1f:9e:a9:c1:40  ucst 966 2 ae2.70

fe80::21f:9eff:fea9:c140/128      dest      0 0:1f:9e:a9:c1:40  ucst 968 2 ae4.90
```

This type of next-hop topology was not seen when Juniper device established an OSPF3 adjacency to another Juniper device. [[PR753849](#): This issue has been resolved.]

- As soon as the forwarding table starts building up on the FPCs, all the affected FPCs will start reporting the following JTREE errors which is an indication of this issue:

```
Apr 29 13:45:54 lab-router fpc5 JTREE(jt_nh_get_reachable_nh32): Not reachable
0x00000000:0x082d1782 for seg 1 (rt_jtree_build_nh)
```

Apr 29 13:45:54 lab-router fpc5 RT: Failed prefix add IPv4 - 1.0.0/24 (jtree nh build failed) on FE 0

Apr 29 13:45:54 lab-router fpc5 RT: IPv4:0 - 1.0.0/24 (add rt entry into jtree failed)

These messages will be seen as soon as the forwarding table starts building up, even if there is no traffic. When traffic starts flowing, the FPCs might crash as a result of this corrupted JTREE. This issue is not seen if per-packet load balancing is configured on the router for all prefixes. [[PR756464](#): This issue has been resolved.]

- When an FPC restart is performed, some of the PICs and physical interfaces are unable to be created by chassisd due to an EBUSY error returned by the kernel. The kernel is unable to process the new requests until the previous states of the same object (PIC, physical interface in our case) are consumed by all peers interested in this. The enhancement addresses the design that makes sure new state changes which could have been processed by the faster peers are not blocked due to these slower peers. [[PR769632](#): This issue has been resolved.]
- Certain hardware data structures used for replicating packets on a single Packet Forwarding Engine stream (SSM list) do not get updated when the corresponding next hops get modified, resulting in use of stale data for multicast replication. All applications that depend on packet replication (IP multicast, P2MP, VPLS BUM traffic) are impacted. Packets are either sent out on wrong logical interfaces or are dropped. However, this happens only if the next hops used for packet replication are modified. This affects line cards using the I/J Chipset in MX, TX, and M Series chassis. [[PR776149](#): This issue has been resolved.]

### ***High Availability (HA) and Resiliency***

- During high routing churn, a flapping interface can in some rare circumstances result in the replicated (backup) kernel to panic with reason "<interface-name>: bitstring index 14 not empty for <mac-address>." [[PR698608](#): This issue has been resolved.]
- If one or more Packet Forwarding Engine peers are slow in consuming ifstates, the secondary Routing Engine does not send CP ACK to the master Routing Engine within a prescribed time. As a result, the secondary Routing Engine is assumed to be having a problem, and hence the connection for the secondary Routing Engine peer is reset, so that ksyncd can clean up the ifstates on the secondary Routing Engine and resync with master Routing Engine again. With this fix, if the secondary CP ACK does not arrive in a prescribed time, if there is any Packet Forwarding Engine that is causing this delay, the same is logged and the CP ACK timer is reset. If no peers are found to be causing the delay of the secondary CP ACK, the behavior is retained to reset the secondary Routing Engine connection. [[PR727344](#): This issue has been resolved.]
- The MPC can generate a core file during unified ISSU. This issue is intermittent. [[PR744992](#): This issue has been resolved.]
- When performing unified ISSU on 10GE DPC, the peer device will see link flap. [[PR777798](#): This issue has been resolved.]
- On MX Series routers with ADPC line cards, performing a unified ISSU upgrade to Junos OS Release 11.4R3.7 might cause the ADPC line cards to reset, thus impacting router operation and defeating the purpose of unified ISSU. Customers with ADPC line cards

on MX Series routers should upgrade only in a maintenance window during which the resets can be tolerated. Other platforms are not affected by this bug, and unified ISSU works as expected. ADPC line cards might reset, causing a traffic outage of approximately 150 seconds duration. [PR779348: This issue has been resolved.]

### **Infrastructure**

- Certain system resources might become exhausted during the Routing Engine switchover under heavy load, causing the system to restart. After restart, the router will operate as expected. [PR733555: This issue has been resolved.]
- A socket hole in the received sequence space on the backup Routing Engine and that backup Routing Engine cannot handle the TCP SACK from the master Routing Engine properly. When backup Routing Engine becomes the new master Routing Engine by switchover, this potential sequence mismatch in the previous backup Routing Engine comes out on the new master Routing Engine. Therefore, this message is generated on syslog. Once after the new master Routing Engine starts handling TCP SACK properly, this mismatch will be cleared and sooner or later this message stops. This is just a cosmetic issue. [PR743382: This issue has been resolved.]
- Fetching ppX interface statistics leaks in pfestat\_table leads to "pfestat\_req\_add: pfestat table out of ids" error logs. When in this state, it is not possible to fetch any interface statistics. To recover from this issue, reload the Routing Engine. Products affected by this are non-MX Series products that offer PPPoE services. [PR751366: This issue has been resolved.]
- Arp entries are not flushed out after disabling interface with purging,aging-timer configured on local router. [PR753268: This issue has been resolved.]
- In scenarios where "family inet" is configured in the pppoe dynamic profile, if ARP request is received on the pppoe interface, the kernel crashes. [PR769646: This issue has been resolved.]

### **Interfaces and Chassis**

- Error message seen on MX80 "fru\_is\_present: out of range slot -1 for CB" continuously. [PR540868: This issue has been resolved.]
- "HS Link FIFO underflow" errors might occur as traffic egresses a PIC when the ingress interface is on another PIC in the same MX-FPC. The speed of the interface and the traffic pattern is relevant to this problem. [PR687905: This issue has been resolved.]
- "ezchip\_xeth\_add\_pw\_bd\_table\_entry" error message is seen on restarting the ADPC card and total drop in traffic is observed after that. This issue will be seen if the following conditions are satisfied:
  - VPLS routing instance is created with configuration "protocol vpls connectivity-type permanent".
  - LSI interface for the vpls routing-instance has a Primary/secondary MPLS LSP present on the ADPC physical interface which is the DUT. Now on just rebooting the ADPC, these logs are seen on the DPC console.

[PR693066: This issue has been resolved.]

- MPC2 might reboot when swapping MIC cards in the same MPC [[PR728095](#): This issue has been resolved.]
- On T Series ES type of FPC, BFD sessions might get flapped when other PIC on the same FPC is brought online. This is caused by the fact that the PIC drivers take long time to do initialization when being brought up which might cause the BFD thread to lose chances of processing the keepalive packets and hence drop the sessions. [[PR733657](#): This issue has been resolved.]
- The issue is present in MX Series platform MIPs must place their own MAC address in the Egress Identifier TLV in CFM Linktrace Messages that they process. They are incorrectly leaving this value unchanged. [[PR735419](#): This issue has been resolved.]
- Due to an incorrect calculation, memory heap utilization of a service PIC can go over 100% under the **show chassis pic** command. There is no service impact. [[PR737676](#): This issue has been resolved.]
- In an Active/Active MC-LAG scenario, traffic might get dropped if:
  - Upstream and downstream interfaces are MC-AE interfaces
  - You have routing protocols running over the IRB
  - Traffic crosses the ICL[[PR746055](#): This issue has been resolved.]
- In Junos OS releases prior to 11.2, the BERT test results would report Error Bit/LOS sec for a newly confirmed E1 link in unframed mode. The issue would be seen on CHSTM1-IQ and CHE1T1-IQE pic. Due to known hardware limitation on CHSTM1-IQ pic, this issue persists on this pic type. However for CHE1T1-IQE pic it has been fixed on Junos OS Releases 11.2R7 and later. [[PR748175](#): This issue has been resolved.]
- If rlsq interfaces are part of a routing instance, upon deactivation and activation of routing instance, all rlsq interfaces were not brought up. This issue is fixed as part of this PR. [[PR749760](#): This issue has been resolved.]
- On a GE port with optic SFP-FX which has auto-negotiation disabled, it might show up even though no cable is connected and have an issue with traffic forwarding on the interface. [[PR751536](#): This issue has been resolved.]
- This issue is specific to I-chip-based DPCs/FPCs and impacts all type of multicast traffic such as IP multicast packets, or L2 multicast/broadcast packets going through L2VPN/VPLS. An I-chip-based DPC/FPC will only forward multicast traffic to the first 1024 receivers of a multicast group if the total number of receivers on a particular PIC of the DPCE, for that group, is between 1025 and 1088 (1024+64). [[PR752662](#): This issue has been resolved.]

### **Layer 2 Features**

- In a router running a VPLS configuration, an administrator configuration change or a network event that causes the removal of an IFF from a VPLS instance could lead to a panic on the backup Routing Engine. [[PR750036](#): This issue has been resolved.]
- The Routing Engine kernel crash was caused by a suspicious packet in the wrong system queue. Packet was classified as a TNP packet (ethertype: 0x8850). TNP is a L3 protocol

used for inter process communication between the Routing Engine and the Packet Forwarding Engine. [[PR779079](#): This issue has been resolved.]

### **Layer 2 Ethernet Services**

- With the configuration of STP/AE under IRB interface, you might see kernel panic on both master/backup Routing Engines after a multiple GRES switchover is done. [[PR742940](#): This issue has been resolved.]
- There is a limitation in the support of IRB interfaces used with DHCP such that:
  - If an LT interface is configured as the underlying interface for an IRB interface and a DHCP client requests a unicast response, then instead of rejecting the send operation a corrupt packet is sent.
  - If the underlying interface of an IRB interface has a different number of tags configured than the bridge domain of the IRB interface and a DHCP client requests a unicast response, a malformed packet is sent.
  - Regardless of tag configuration on the bridge domain, if the packet needs to be relayed out a VPLS tunnel (an LSI or VT interface as underlying), a malformed packet is sent.

[[PR751398](#): This issue has been resolved.]

### **MPLS**

- By design, family MPLS under virtual-router type routing-instance does not get created without a corresponding "protocols:ldp". Hence, without MPLS family, MPLS filter is not working on an interface configured under virtual-router type routing-instance. As a workaround, configure ldp in the instance, and disable it on all interfaces if not used. [[PR601989](#): This issue has been resolved.]
- The routing protocol process might redundantly try to save the RSVP ERO object in the graceful restart database. This is applicable only for non-traffic engineered LSPs when graceful restart is configured. [[PR741694](#): This issue has been resolved.]
- l3vpn-composite-nexthop with MPC and MSDPC doing stateful-firewall with interface service-set will drop packets on service input direction. The interface where service input/output is configured has to be inside a VRF, and the destination for which service input should intercept the traffic and send it to service PIC in MSDPC should be reachable through MPLS backbone, so resolved through composite next-hop, in order to see this issue. [[PR747914](#): This issue has been resolved.]
- Traffic fails to go through service output when it comes from MPLS core and is routed inside VRF without vrf-table-label configured. This should NOT work on all types of FPCs except MPCs on MX Series routers. This PR fixes the problem on MPCs. [[PR749661](#): This issue has been resolved.]
- When LSP is configured with auto-bandwidth switches from the primary path to secondary path, bandwidth estimation on the secondary path might be under-estimated. Due to under-estimation, overflow sample count might get reset. [[PR752777](#): This issue has been resolved.]

- The kernel might crash at `tag_mtu_calc` when the Routing Engine attempts to send a packet larger than the configured MPLS MTU, warranting fragmentation (over a LSP) using a `l3vpn-composite-nexthop`. For the issue to occur both must be true: 1) `l3-composite-nexthop` knob must be turned on. 2) MPLS MTU must be manually configured by the user. [[PR755950](#): This issue has been resolved.]
- On MX Series router with MICs/MPCs, when switch L2 MPLS packets on egress PE routers, the inner MPLS label TTL value is checked and if valid decreased by 1. During the egress process, the TTL value is rechecked. If the value is 1 at this point, the packet is sent to the Routing Engine instead of being forwarded out the interface. [[PR776203](#): This issue has been resolved.]

### ***Platform and Infrastructure***

- Packets exchanged between logical routers within the same physical router over logical tunnel (LT) interfaces will not have their TTL decremented. [[PR685639](#): This issue has been resolved.]
- Under very special race conditions, the MPC CPU might stop processing and will be reset due to Level3/Level 2 watchdog expiration timer. Potential exposure causes high load of traffic sent to the Host. The following syslog message will be reported in the syslog once MPC reboots. "fpc[x] MPC: Reset reason (0xc): Level3 watchdog, Level2 watchdog." [[PR717899](#): This issue has been resolved.]
- Enabling of Dynamic Profile versioning is not supported if dynamic profiles have been already configured on the router. If you deactivate existing dynamic profiles in order to enable and commit dynamic profile versioning, profile version numbers are not subsequently incremented. As a workaround, you must delete all existing dynamic profiles before you enable profile versioning, and then reconfigure the dynamic profiles. [[PR741001](#): This issue has been resolved.]
- This is an issue during the passing of timestamp message from kernel to rmopd for 64-bit Junos OS. [[PR746428](#): This issue has been resolved.]
- If a filter contains multiple prefix actions and the filter is applied, changing one prefix action referenced by this filter might crash NPCs. The change on a prefix action could be direct or indirect (for example, changing the policer reference by this prefix action). The workaround is detaching all such filters before changing a prefix action and then applying the filters back after the change. [[PR750370](#): This issue has been resolved.]
- When CoS rewrite is configured for an IRB interface, and the IRB interface participates in L2 multicast, the copies sent over the physical interface will not have the CoS rewrite applied. This issue is applicable only when the chassis is configured in the "enhanced-ip" mode. [[PR754720](#): This issue has been resolved.]
- A MPC-\* FPC installed in an MX240/480/960 router or the integrated TFEB of an MX5/10/40/80 router might crash and reboot when the unsupported command **show route hw nhs** is executed from the FPC CLI. This command is unsupported and should not be used without the explicit instructions of JTAC. It is not needed for the normal operation of a Juniper Networks router. [[PR772413](#): This issue has been resolved.]
- If "source-filtering" is turned on under an interface, packets with a multicast destination mac address will get dropped. Such packets are used by applications like CFM. The



multicast mac addresses cannot be explicitly added to be accepted using the CLI. [PR772611: This issue has been resolved.]

- Traffic-control-profile applied on LT logical interface used to terminate a vpls instance has no effect and the logical interface is not shaped. [PR773764: This issue has been resolved.]
- Customers using Junos OS Release 11.4R3.6 code on MX Series routers with MPC 3D 16x 10GE line cards might experience issues with interfaces on these line cards. Some interfaces on the MPC 3D 16x 10GE line cards might be reported as UP ("Enabled" and "Physical Link UP") in the **show interfaces <interface>** command. However, **show interfaces <interface> terse** command for the same interface reports that interface as DOWN (Admin - UP and Link Protocol - DOWN). The link lights at both ends of the link will be GREEN, thereby indicating connectivity. However, no traffic passes through the affected interfaces. This issue was seen on interfaces that were part of aggregated Ethernet (AE) bundle as well as on interfaces that were not part of the AE bundle. In addition, "Wedge Detected" messages might be seen in the syslogs and in the telnet/ssh session to the router. This behavior was not seen with DPC hardware. [PR776727: This issue has been resolved.]

### ***Routing Protocols***

- The **show bgp group** output is updated to new multiline format in order to display the full name of table bgp.rtarget.0. [PR696476: This issue has been resolved.]
- ISO/CLNS prefixes with more than /152 VPN prefix length when advertised by BGP across VPN core causes BGP adjacent flap since the remote BGP rejects the same prefix as an invalid address. This is because the ISOVPN draft allows only up to /152 prefixes. [PR742491: This issue has been resolved.]
- Pruned multicast traffic continues to flow from the source even when receiver leaves the multicast group for Junos OS Releases 10.4R8.5, 10.4R9, and 10.4R9.2. [PR746474: This issue has been resolved.]
- Route advertisement stops for RT family enabled BGP peers after VRF is deactivated and activated. This issue is only seen with RT-enabled peers and nonstop active routing enabled. [PR749288: This issue has been resolved.]
- If there is some link micro flapping, it might bring the BFD into a problematic state. As a result, for next event of BFD state down, it will not bring down the client sessions like OSPF, IS-IS, BGP, etc. [PR749388: This issue has been resolved.]
- The routing protocol process crashes and generates a core file after executing **show ospf context-identifier area <area>** command which is given for an area that has not been configured. The issue is caused by insufficient check code. [PR750914: This issue has been resolved.]
- The multipath flash mechanism runs unnecessarily when BGP multipath is configured for inet-vpn routes. When large amounts of inet-vpn routes change, there is a noticeable delay in convergence for the inet-vpn routes. [PR751469: This issue has been resolved.]
- If BGP receives an ISO-VPN prefix of length 248, i.e. ISO part of prefix contains NSEL-byte, BGP session will be reset. This is according to standards, but it would be good if BGP can handle this gracefully without resetting BGP-session. This PR makes

BGP handle it gracefully, by ignoring the NSEL byte received in the ISO-VPN prefix. [PR771835: This issue has been resolved.]

- Customer needs these debug logs to be changed to severity LOG\_DEBUG. "mcsn[91713]: krt\_decode\_nexthop: Try freeing: nh-handle: 0x0 nh-index: 1049040 fwdtype: 2" This was introduced as part of Release 11.4 with severity set to "LOG\_INFO" (will not be seen with earlier releases). This is used as a debug log and is harmless. "mcsn[91713]: Received MC\_AE\_OPTIONS TLV for intf device ae1; mc\_ae\_id 0, status 2" This was introduced as part of RLI 8857 in Release 10.0 (reference from PR-411614). This also has a severity of "LOG\_INFO" and is part of the rpd-infra that is used by mcsnoopd. [PR772063: This issue has been resolved.]
- The routing protocol process might generate a core file while processing malformed RIP or RIPng message from neighbor during adjacency establishment. [PR772601: This issue has been resolved.]
- Limited Support for multiple area TLVs in a single IS-IS Hello message: When many area TLVs are found in a single IS-IS Hello packet, L1 adjacencies might not be formed correctly and can be stuck in the initializing state. Currently, there are no identified workarounds; however, this does not impact L2 adjacencies. [PR775852: This issue has been resolved.]

### **Services Applications**

- When you pump in more than 2.1 million passive monitoring flows into Monitor-II PIC, the router might not send memory overload SNMP trap. [PR677162: This issue has been resolved.]
- When sending traffic through IPsec tunnels for above 2.5Gbps on an MS-400 PIC, the Service-PIC might bounce due to prolonged flow control. [PR705201: This issue has been resolved.]
- If the Service PIC processing DS-Lite packets receive packets from overlapping IPv4 addresses present behind different B4s at the same time, then there is a possibility that the PIC will crash with a similar cored file. [PR711307: This issue has been resolved.]
- "linerate-mode" might not be applied correctly to interfaces when first configured on a PIC which does not support "linerate-mode" and later on replace the first PIC with a second PIC which supports "linerate-mode". [PR734887: This issue has been resolved.]
- In Junos OS Releases 10.4 and later, the number of outstanding IPsec tunnels has changed to be 50 tunnels instead of 200 outstanding tunnels in previous releases. [PR739683: This issue has been resolved.]
- DCD\_CONFIG\_WRITE\_FAILED with "Device not configured" error is observed when system is rebooted with rlsq configuration. As a workaround, deactivate interfaces, request a system reboot, and activate interfaces, instead of only requesting a system reboot. [PR741121: This issue has been resolved.]
- This PR enables visibility of Address Pool Paired out of port errors using the CLI command **show services nat pool detail**.

```
user@router-re0> show services nat pool detail
Interface: sp-7/0/0, Service set: nat44
NAT pool: public-pool, Translation type: dynamic
```

Address range: 100.100.0.1-100.100.0.254  
 Port range: 512-65535, Ports in use: 64512, Out of port errors: 0, Max ports used: 64512  
 AP-P out of port errors: 440601 <-- errors are now shown here

[[PR746752](#): This issue has been resolved.]

- This is a memory leak in the IDPD daemon on the Routing Engine. It occurs when SNMP queries are done on the Routing Engine. This leak is relatively slow and occurs over several days. When the size of the daemon reaches 512M, it generates a core file.  
 [[PR748414](#): This issue has been resolved.]

- There is no logical binding of <flow-analysis-statistics-entry> to <flow-analysis-statistics-pic-info> in the output of **show services stateful-firewall flow-analysis | display xml**. At the moment pairs of these tags are just put sequentially on the same level of hierarchy under <service-flow-analysis-information> as can be seen in the following output:

```
user@router> show services stateful-firewall flow-analysis | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/10.458/junos">
  <service-flow-analysis-information>
    <flow-analysis-statistics-pic-info>
      <pic-name>sp-0/0/0</pic-name>
    </flow-analysis-statistics-pic-info>
    <flow-analysis-statistics-entry>
      .... output omitted ....
    </flow-analysis-statistics-entry>
    <flow-analysis-statistics-pic-info>
      <pic-name>sp-0/1/0</pic-name>
    </flow-analysis-statistics-pic-info>
    <flow-analysis-statistics-entry>
      .... output omitted ....
    </flow-analysis-statistics-entry>
  ... output omitted ...
```

The Junos OS software has been modified to include a new tag <service-flow-analysis-entry>, which is the parent of both the <flow-analysis-statistics-entry> tag and the <flow-analysis-statistics-pic-info> tag, thus tying the pic name with existing flow analysis details:

```
user@router> show services stateful-firewall flow-analysis | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/12.1I0/junos">
  <service-flow-analysis-information>
    <service-flow-analysis-entry>          <----- Start New
tag
    <flow-analysis-statistics-pic-info>
      <pic-name>sp-2/0/0</pic-name>
    </flow-analysis-statistics-pic-info>
    <flow-analysis-statistics-entry>
      .... output omitted ....
    </flow-analysis-statistics-entry>
    <flow-analysis-num-flows-sec-samples-entry>
      .... output omitted ....
    </flow-analysis-num-flows-sec-samples-entry>
    <flow-analysis-num-flows-sec-entry>
      .... output omitted ....
    </flow-analysis-num-flows-sec-entry>
    <flow-analysis-num-flows-sec-entry>
      .... output omitted ....
```

```

        </flow-analysis-num-flows-sec-entry>
        <flow-analysis-protocol-lifetime-entry>
        .... output omitted ....
        </flow-analysis-protocol-lifetime-entry>
    </service-flow-analysis-entry>          <----- End
    <service-flow-analysis-entry>          <----- Start New tag

        <flow-analysis-statistics-pic-info>
        <pic-name>sp-2/1/0</pic-name>
    </flow-analysis-statistics-pic-info>
    <flow-analysis-statistics-entry>
    .... output omitted ....
    </flow-analysis-statistics-entry>
    <flow-analysis-num-flows-sec-samples-entry>
    .... output omitted ....
    </flow-analysis-num-flows-sec-samples-entry>
    <flow-analysis-num-flows-sec-entry>
    .... output omitted ....
    </flow-analysis-num-flows-sec-entry>
    <flow-analysis-num-flows-sec-entry>
    .... output omitted ....
    </flow-analysis-num-flows-sec-entry>
    <flow-analysis-protocol-lifetime-entry>
    .... output omitted ....
    </flow-analysis-protocol-lifetime-entry>
    </service-flow-analysis-entry>          <----- End
</service-flow-analysis-information>
<cli>
    <banner>[edit]</banner>
</cli>
</rpc-reply>

```

[[PR749675](#): This issue has been resolved.]

- While trying to Allocate NAT ports for SIP headers, MS-PIC crashes. [[PR769605](#): This issue has been resolved.]

### **User Interface and Configuration**

- In edit private mode, when a node is disabled and then annotate is used on disable, it throws error on commit. [[PR58358](#): This issue has been resolved.]
- Using the "# load" command to replace policy configuration could lead to a configuration corruption which causes rpd to crash upon commit. [[PR704294](#): This issue has been resolved.]

### **VPNs**

- An optimization has been implemented with BGP-MVPN next-hop infrastructure which will improve scalability in some multi-dimensional scaling scenarios with aggregate interfaces. [[PR690690](#): This issue has been resolved.]
- Under certain circumstances a vrf-import policy's term with the "accept" action that matches the BGP VPN route based on the criteria different than the target community can reject the matching route. [[PR706064](#): This issue has been resolved.]
- Currently, MVPN Leaf-AD routes with IR provider tunnels are sent without the PMSI attributes. These routes should be sent with the PMSI attributes. The label will be the same label as advertised in the Type 1 route. [[PR717451](#): This issue has been resolved.]

- In an MVPN scenario, performing GRES might result in some traffic loss for MVPN flows. [[PR733893](#)]
- With BGP MVPNs when there are many interfaces in the vrf, it is possible that RPD might core. If a forwarding entry has a large number of outgoing interfaces, this memory error will occur. The exact number of outgoing interfaces needed to trigger this issue is not known. [[PR749379](#): This issue has been resolved.]
- In NG-MVPN with a multihomed source attached to ingress PE, when the original-DR goes down and then comes back to claim its role as DR, the other node will lose its intermediate DR-role and withdraw its type 5 AD-route. However, the new DR that comes back will not advertise a type 5 AD route. As result of this misbehavior, neither the non-DR nor the DR will advertise a type 5 AD route in the re-convergence case and hence no egress-PE could join the source. [[PR754222](#): This issue has been resolved.]
- The issue happens when the ingress PE receives the type-4 leaf AD route before discovering the egress PE as a neighbor using a type-1 route. PE ignores the type-4 leaf AD route as there is no nbr. When the ingress PE receives the type-1 route, it only processes inclusive p-tnl and since it did not add the unicast IR tunnel as a leaf to the spmsi tunnel, the egress PE doesn't receive the traffic. [[PR755209](#): This issue has been resolved.]
- When the label for intra-AS AD route changes, it is not reflected in the intra-as AD route generated to the MVPN PE peers. As a result the peers still use the old label information and results in traffic drop. [[PR771059](#): This issue has been resolved.]

**Related Documentation**

- [New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 100](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 154](#)
- [Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 170](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 285](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 322](#)

## Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers

### Errata

---

- [Hardware](#)
- [Class of Service \(CoS\)](#)
- [High Availability \(HA\) and Resiliency](#)
- [Interfaces and Chassis](#)
- [Infrastructure](#)

- [Junos OS XML API and Scripting](#)
- [J-Web Interface](#)
- [Layer 2 Ethernet Services](#)
- [Multicast](#)
- [MPLS](#)
- [Network Management](#)
- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Services Applications](#)
- [Subscriber Access Management](#)
- [Timing and Synchronization](#)
- [User Interface and Configuration](#)
- [VPNs](#)

#### **Hardware**

- The *Protocols and Applications Supported by MX240, MX480, MX960, MX2010, and MX2020 MPCs* topic erroneously states that support was introduced in Junos OS Release 10.4 for IEEE 802.3ah OAM (discovery and link monitoring, fault signaling and detection, and remote loopback). In fact, this support was introduced in Junos OS Release 11.1.

#### **Class of Service (CoS)**

- The *Example: Configuring Scheduling Modes on Aggregated Interfaces* topic fails to mention the following additional information regarding the parameters that are scaled for aggregated interface member links when the scheduler parameters are configured using scheduler maps:

Apart from transmit rate and buffer size that are scaled when the parameters are configured using scheduler maps, shaping rate is also scaled if you configure it in bits per second (bps). Shaping rate is not scaled if you configure it as a percentage of the available interface bandwidth.

[*Class of Service, Schedulers on Aggregated Ethernet and SONET/SDH Interfaces*]

- The following additional information regarding the processing of custom EXP rewrite rules on MPCs applies to the *Rewriting IEEE 802.1p Packet Headers with an MPLS EXP Value* topic:

For MPCs, default EXP rewrite rules do not exist for logical interfaces. The EXP CoS bits for MPLS labels are obtained from the IP precedence bits for IP traffic. The EXP bits for labels that are pushed or swapped are inherited from the current label of the MPLS packets. For non-IP and non-MPLS packets, the EXP bits are set to 0. If a custom EXP rewrite rule is configured on the core-facing interface, then it overrides the EXP bits.

[*Class of Service, CoS Re-Marking of Packets Entering or Exiting the Network*]

**High Availability (HA) and Resiliency**

- 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]

- The MX Series Virtual Chassis documentation in the *Junos OS High Availability Configuration Guide* failed to include the following information about how slot numbering in the Virtual Chassis affects your use of SNMP.

Junos OS supports the use of SNMP to monitor the routers and other devices in your network. For example, the Juniper Networks jnxBoxAnatomy enterprise-specific Chassis MIB contains the jnxFruTable object, which shows the status of field-replaceable units (FRUs) in the chassis. Within the jnxFruTable object, the jnxFruSlot object displays the slot number where the FRU is installed.

If you are using the jnxFruSlot object in jnxFruTable to display the slot numbers of line cards installed in a member router of an MX Series Virtual Chassis, keep in mind that the offset used for slot numbering in an MX Series Virtual Chassis affects the value that appears for the jnxFruSlot object.

[Table 4 on page 287](#) lists the jnxFruSlot number that appears in the jnxFruTable of the jnxBoxAnatomy MIB, and the corresponding line card physical slot number in each member router of a two-member MX Series Virtual Chassis. For example, a jnxFruSlot value of 15 corresponds to physical slot 3 in member 0 of an MX Series Virtual Chassis. A jnxFruSlot value of 30 corresponds to physical slot 6 in member 1 of an MX Series Virtual Chassis.

**Table 4: jnxFruSlot Numbers and Corresponding Slot Numbers in an MX Series Virtual Chassis**

jnxFruSlot Number	Line Card Slot Number	MX Series Virtual Chassis Member ID
Line Cards in MX Series Virtual Chassis Member ID 0 (offset = 12):		
12	0	0
13	1	0
14	2	0
15	3	0
16	4	0
17	5	0
18	6	0
19	7	0
20	8	0

Table 4: jnxFruSlot Numbers and Corresponding Slot Numbers in an MX Series Virtual Chassis (*continued*)

jnxFruSlot Number	Line Card Slot Number	MX Series Virtual Chassis Member ID
21	9	0
22	10	0
23	11	0
Line Cards in MX Series Virtual Chassis Member ID 1 (offset = 24)		
24	0	1
25	1	1
26	2	1
27	3	1
28	4	1
29	5	1
30	6	1
31	7	1
32	8	1
33	9	1
34	10	1
35	11	1

[*Junos OS High Availability Configuration Guide, Junos OS SNMP MIBs and Traps Reference*]

- In Junos OS Release 11.4 and later releases, the *Example: Replacing a Routing Engine in a Virtual Chassis Configuration for MX Series 3D Universal Edge Routers* topic in the *MX Series Interchassis Redundancy Using Virtual Chassis* pathway page failed to mention that for a replacement Routing Engine shipped from the factory that you plan to install in an MX Series Virtual Chassis member router, you must modify the default factory configuration to enable proper operation of the Virtual Chassis. The documentation has been updated to include this information in Junos OS Release 13.2 and later releases, as follows:



A Routing Engine shipped from the factory is loaded with a default factory configuration that includes the following stanza at the `[edit]` hierarchy level:

```
[edit]
system {
  commit {
    factory-settings {
      reset-virtual-chassis-configuration;
    }
  }
}
```

When this configuration stanza is present, the Routing Engine can operate only in a standalone chassis and *not* in an MX Series Virtual Chassis member router. As a result, if you install this Routing Engine in the standby slot of a Virtual Chassis member router (**member1-re1** in this example), the Routing Engine does not automatically synchronize with the master Routing Engine and boot in Virtual Chassis mode.

To ensure that the standby factory Routing Engine successfully synchronizes with the master Routing Engine, you must remove this standalone chassis configuration stanza from the standby factory Routing Engine and verify that it reboots in Virtual Chassis mode before you install the Junos OS release.

To modify the Routing Engine factory configuration to ensure proper operation of the MX Series Virtual Chassis:

1. Log in to the console of the new Routing Engine as the user **root** with no password.
2. Configure a plain-text password for the **root** (superuser) login.

```
{local:member1-re1}[edit system]
root# set root-authentication plain-text-password
New password: type password here
Retype new password: retry password here
```

3. Delete the standalone chassis configuration.

```
{local:member1-re1}[edit]
root# delete system commit factory-settings reset-virtual-chassis-configuration
```

4. Commit the configuration.

The new Routing Engine synchronizes the Virtual Chassis member ID with the master Routing Engine and boots in Virtual Chassis mode.

5. Verify that the new Routing Engine is in Virtual Chassis mode.

During the boot process, the router displays the following output to indicate that it has synchronized the Virtual Chassis member ID (1) with the master Routing Engine and is in Virtual Chassis mode.

```
...
virtual chassis member-id = 1
virtual chassis mode      = 1
...
```

- For a two-member MX Series Virtual Chassis to function properly, you must enable enhanced IP network services on both member routers when you first set up the Virtual Chassis. If necessary, you can also enable enhanced IP network services for an existing Virtual Chassis.

Enhanced IP network services defines how the router recognizes and uses certain modules. When you set each member router's network services to **enhanced-ip**, only MPC/MIC modules and MS-DPC modules are powered on in the router. Non-service DPCs do not work with enhanced IP network services.

In Junos OS Release 11.4 and later releases prior to Release 13.2, the documentation for MX Series Virtual Chassis fails to mention the required procedures for enabling enhanced IP network services.

Use the following procedure to enable enhanced IP network services as part of the initial Virtual Chassis configuration. Perform these steps immediately after you create the preprovisioned member configuration on the master router, and before you enable graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) on both member routers.

To enable enhanced IP network services when you first set up an MX Series Virtual Chassis:

1. Configure enhanced IP network services on member 0.

- a. Log in to the console on member 0.
- b. Access the chassis hierarchy.

```
[edit]  
user@hostA# edit chassis
```

- c. Configure enhanced IP network services for member 0.

```
[edit chassis]  
user@hostA# set network-services enhanced-ip
```

- d. Commit the configuration on member 0 by using the **commit synchronize** command.



**NOTE:** Immediately after you commit the configuration, the software prompts you to reboot the router. You can proceed without rebooting the router at this point because a reboot occurs when you configure the member IDs to enable Virtual Chassis mode.

---

2. Configure enhanced IP network services on member 1.

- a. Log in to the console on member 1.
- b. Access the chassis hierarchy.

```
[edit]  
user@hostB# edit chassis
```

- c. Configure enhanced IP network services for member 1.

```
[edit chassis]
user@hostB# set network-services enhanced-ip
```

- d. Commit the configuration on member 1 by using the **commit synchronize** command.



**NOTE:** Immediately after you commit the configuration, the software prompts you to reboot the router. You can proceed without rebooting the router at this point because a reboot occurs when you configure the member IDs to enable Virtual Chassis mode.

3. (Optional) After the Virtual Chassis forms, verify that enhanced IP network services has been properly configured.

- a. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis master router (member0-re0).

```
{master:member0-re0}
user@hostA> show chassis network services
```

Network Services Mode: Enhanced-IP

- b. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis backup router (member1-re0).

```
{backup:member1-re0}
user@hostB> show chassis network services
```

Network Services Mode: Enhanced-IP

Use the following procedure to enable enhanced IP network services for an existing Virtual Chassis configuration.

To configure enhanced IP network services for an existing Virtual Chassis:

1. Log in to the console for the master Routing Engine in the Virtual Chassis master router (member0-re0).

2. Access the chassis hierarchy.

```
{master:member0-re0}[edit]
user@hostA# edit chassis
```

3. Configure enhanced IP network services on member 0.

```
{master:member0-re0}[edit chassis]
user@hostA# set network-services enhanced-ip
```

4. Commit the configuration by using the **commit synchronize** command.
5. When prompted to do so, reboot both Routing Engines in each member router forming the Virtual Chassis.

- For Junos OS Releases 11.4, 12.1, 12.2, 12.3R1, and 12.3R2:

```
{master:member0-re0}
user@hostA> request system reboot member 0 other-routing-engine
```

```
user@hostA> request system reboot member 1 other-routing-engine
user@hostA> request system reboot
```

- For Junos OS Release 12.3R3 and later releases:

```
{master:member0-re0}
user@hostA> request system reboot
```

Rebooting all Routing Engines in the Virtual Chassis propagates the enhanced IP network services configuration to both member routers.

6. (Optional) Verify that enhanced IP network services has been properly configured for the Virtual Chassis.

- a. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis master router (member0-re0).

```
{master:member0-re0}
user@hostA> show chassis network services
```

Network Services Mode: Enhanced-IP

- b. Verify that enhanced IP network services is configured on the master Routing Engine in the Virtual Chassis backup router (member1-re0).

```
{backup:member1-re0}
user@hostB> show chassis network services
```

Network Services Mode: Enhanced-IP

- The following additional information applies to the *Virtual Chassis Components Overview* topic in the *Interchassis Redundancy Using Virtual Chassis Feature Guide for MX Series Routers* for Junos OS Release 11.2 and later releases.

When you configure chassis properties for MPCs installed in a member router in an MX Series Virtual Chassis, keep the following points in mind:

- Statements included at the **[edit chassis member *member-id* fpc slot *slot-number*]** hierarchy level apply to the MPC (FPC) in the specified slot number only on the specified member router in the Virtual Chassis.

For example, if you issue the **set chassis member 0 fpc slot 1 power off** statement, only the MPC installed in slot 1 of member ID 0 in the Virtual Chassis is powered off.

- Statements included at the **[edit chassis fpc slot *slot-number*]** hierarchy level apply to the MPCs (FPCs) in the specified slot number on *each* member router in the Virtual Chassis.

For example, if you issue the **set chassis fpc slot 1 power off** statement in a two-member MX Series Virtual Chassis, both the MPC installed in slot 1 of member ID 0 *and* the MPC installed in slot 1 of member ID 1 are powered off.



**BEST PRACTICE:** To ensure that the statement you use to configure MPC chassis properties in a Virtual Chassis applies to the intended member router and MPC, we recommend that you always include the **member**

*member-ID* option before the **fpc** keyword, where *member-id* is 0 or 1 for a two-member MX Series Virtual Chassis.

### *Interfaces and Chassis*

- **SONET/SDH support on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (MIC-3D-4COC3-1COC12-CE)**—This feature is supported in Junos OS Release 12.2R1. However, the documentation for this feature is not available in this release. Documentation for this feature is planned for an upcoming release.
- With Junos OS Release 10.1 and later, you need not include the **tunnel** option or the **clear-dont-fragment-bit** statement when configuring **allow-fragmentation** on a tunnel.

[*Services Interfaces*]

- Hybrid mode is currently not supported in Junos OS Release 12.2R1. All references to hybrid mode (combined operation of Precision Time Protocol and Synchronous Ethernet) in the *System Basics Configuration Guide* and *Junos System Basics and Services Command Reference Guide* should be disregarded.
- The **show chassis fabric reachability** and the **show chassis fabric unreachable-destinations** command topics fail to state that these commands are also supported on MX240, MX480, and MX960 routers from Junos OS Release 11.4R2 and Junos OS Release 12.1. The Supported Platforms section of this topic fails to mention MX240, MX480, and MX960 routers on which these commands are supported.

[*System Basics and Services Command Reference*]

- The *Interfaces and Chassis* subsection in the *New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers* section of the Junos OS 12.2R1 Release Notes incorrectly contains the *Support for disabling an FPC with degraded fabric bandwidth* topic. This topic applies to the Junos OS 12.1 and Junos 11.4 Release Notes because this feature is introduced in Junos OS Release 11.4R3 and Junos OS Release 12.1R1.

[*Release Notes*]

- The *Interfaces and Chassis* subsection in the *New Features in Junos OS Release for M Series, MX Series, and T Series Routers* section of the Junos OS 12.2R1 and Junos OS 11.4R3 Release Notes fails to describe the following information regarding support for redundancy fabric mode for active control boards on MX Series routers. This feature is available in Junos OS Release 11.4R3 and later, and Junos OS Release 12.2 and later.

To configure redundancy mode for the active control board, use the **redundancy-mode redundant** statement at the [**edit chassis fabric**] hierarchy level. When you configure this option, all the FPCs use 4 fabric planes as active planes, regardless of the type of the FPC. If you do not configure this option, increased fabric bandwidth mode is enabled by default on MX Series routers with Switch Control Board (SCB). For the MX Series routers that contain the enhanced Switch Control Board (SCB) with Trio chips and the MPC3E, the control boards operate in redundancy fabric mode (all the FPCs use 4 fabric planes as active planes) by default.

To configure increased bandwidth mode for the active control board, use the **increased-bandwidth** statement at the **[edit chassis fabric]** hierarchy level. When you configure this option, all the available fabric planes are used.

Configuring this feature does not affect the system. You can configure this feature without restarting the FPC or restarting the system.

You can use the **show chassis fabric redundancy-mode** command to verify whether the redundancy fabric mode is enabled.

[Release Notes]

- **Support for redundant fabric made on active control boards of MX Series routers**—An FPC working with reduced fabric bandwidth can affect the re-routing process and can cause partial traffic black holes. You can now enable increased fabric bandwidth of active control boards for optimal and efficient performance and traffic handling. On an MX960, MX480, or MX240 router, you can configure the active control board to be in redundancy mode or in increased fabric bandwidth mode. In increased fabric bandwidth mode, the maximum number of available fabric planes are used for MX Series routers with Trio chips and the MPC3E. On MX960 routers with active control boards, 6 active planes are used, and on MX240 and MX480 routers with active control boards, 8 active planes are used.
- The **show chassis fabric destinations fpc <fpc-slot-number>** command is supported on MX240, MX480, and MX940 routers in Junos OS Release 12.1 and later. This command can be used to display the state of fabric destinations for all FPCs or a particular FPC.

[Table 5 on page 294](#) lists the output fields for the **show chassis fabric destinations** command. Output fields are listed in the approximate order in which they appear.

**Table 5: show chassis fabric destinations Output Fields**

Field Name	Field Description
<b>Fabric destinations state</b>	Indicates the state of the fabric destinations: <ul style="list-style-type: none"> <li>• <b>0</b>—Destination is non-existent.</li> <li>• <b>2</b>—Destination is enabled.</li> <li>• <b>3</b>—Destination is disabled.</li> <li>• <b>6</b>—Destination is in erroneous state and is disabled.</li> </ul>
<b>Flexible PIC Concentrator (FPC) number</b>	Source FPC number.
<b>Packet Forwarding Engine number</b>	Source Packet Forwarding Engine number.
<b>Plane number</b>	Source plane number.

The following sample output displays the state of fabric destinations for all FPCs on an MX Series router:

```
user@host> show chassis fabric destinations fpc 1
```

```
Fabric destinations state:
```

0: non-existent  
2: enabled  
3: disabled  
6: dest-err and disabled

```
FPC 1
PFE 0
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
PFE 1
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
```

The following sample output displays the state of fabric destinations for a particular FPC on an MX Series router:

```
user@host> show chassis fabric destinations
```

```
Fabric destinations state:
0: non-existent
2: enabled
3: disabled
6: dest-err and disabled
```

```
FPC 1
PFE 0
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
PFE 1
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
FPC 2
PFE 0
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
```

```
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
PFE 1
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
PFE 2
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
PFE 3
Plane 0  0000 3300 3333
Plane 1  0000 2200 2222
Plane 2  0000 2200 2222
Plane 3  0000 2200 2222
Plane 4  0000 2200 2222
Plane 5  0000 3300 3333
Plane 6  0000 3300 3333
Plane 7  0000 3300 3333
```

- The **tunnel-services** configuration statement topic incorrectly states that you can use the **tunnel-services** statement to specify that the IQ2 or IQ2E PIC will work both as a regular PIC and as a tunnel PIC. The correct functionality of the **tunnel-services** statement is as follows:

You can specify the IQ2 and IQ2E PICs to work exclusively in tunnel mode or as a regular PIC. To configure exclusive tunnel mode, use the **tunnel-only** statement at the **[edit chassis fpc slot-number pic slot-number tunnel-services]** hierarchy level. The default setting uses IQ2 and IQ2E PICs as a regular PIC. If you do not configure the **tunnel-only** option, the IQ2 and IQ2 PICs operate as regular PICs.

*[System Basics, Chassis-Level Features]*

- The **frame-error** configuration statement topic incorrectly states that the default window during which frame errors are counted until they reach the configured threshold is 100 milliseconds. The correct description of the default window is as follows:

The window or period during which frame errors are counted is 5 seconds or multiples of it (with a maximum value of 1 minute). This window denotes the duration as intervals of 100 milliseconds, encoded as a 16-bit unsigned integer. This window is not configurable in Junos OS. According to the IEEE 802.3ah standard, the default value of the frame-errors window is 1 second. This window has a lower bound of 1 second and an upper bound of 1 minute.

*[Network Interfaces, Ethernet Interfaces]*



- In the *Chassis Conditions That Trigger Alarms* section, the following additional information regarding the generation of alarms when the management interface is down in routers with a single Routing Engine or the master Routing Engine applies to Table 1 through Table 8.

**Table 6: Chassis Component Alarm Conditions**

Chassis Component	Alarm Condition	Remedy	Alarm Severity
Routing Engine	The Ethernet management interface ( <b>fxp0</b> or <b>em0</b> ) on the Routing Engine is down.	<ul style="list-style-type: none"> <li>• Check the interface cable connection.</li> <li>• Reboot the system.</li> <li>• If the alarm recurs, open a support case using the Case Manager link at <a href="http://www.juniper.net/support/">http://www.juniper.net/support/</a> or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).</li> </ul>	Red

[*System Basics, Chassis-Level Features*]

- The **lmi-type** statement incorrectly states that Consortium LMI is supported only on M320 routers with Enhanced III FPCs and specific IQE PICs and on MX80, MX240, MX480, and MX960 routers with MICs specified in the *Configuring Tunable Keepalives for Frame Relay LMI* section. The following is the correct compatibility statement:

Consortium LMI is supported on all MPCs and I-chip based FPCs.

[*Frame Relay Interfaces*]

- The *Support for redundant fabric made on active control boards of MX Series routers* topic in the *Interfaces and Chassis* subsection in the *New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers* section of the Junos OS 12.2R2 Release Notes and Junos OS 11.4R6 Release Notes erroneously states that If you do not configure the **redundancy-mode redundant** statement at the **[edit chassis fabric]** hierarchy level, increased fabric bandwidth mode is enabled by default on MX Series routers.

The correct default behavior of redundant fabric mode on MX Series routers is as follows:

Increased fabric bandwidth mode is enabled by default on MX Series routers with Switch Control Board (SCB). On MX Series routers that contain the enhanced SCB with Trio chips and the MPC3E, redundancy mode is enabled by default.

[*Release Notes*]

- The **forwarding-mode (100-Gigabit Ethernet)** configuration statement topic fails to mention that this statement is supported on MX Series routers from Junos OS Release 12.1. The Supported Platforms section of this topic fails to list MX Series routers on which this command is supported.

[*Network Interfaces, Ethernet Interfaces*]

- The *IP Demux Interfaces over Static or Dynamic VLAN Demux Interfaces* topic incorrectly states that both DPCs and MPCs support VLAN demux subscriber interfaces. In fact, only MPCs support these interfaces.
- The **show chassis fabric unreachable-destinations** command is incorrectly mentioned as supported on MX240, MX480, and MX960 routers from Junos OS Release 11.4R2 and Junos OS Release 12.1. The Supported Platforms section of this topic also incorrectly states that MX240, MX480, and MX960 routers as supported routers for this command. This command is not available on the MX240, MX480, and MX960 routers. Instead, the correct command is the **show chassis fabric destinations** command, which you can use to view the state of fabric destinations for all FPCs.

[*System Basics and Services Command Reference*]

- The *Limiting traffic black-hole time by detecting Packet Forwarding Engine destinations that are unreachable over the fabric (MX240, MX480, and MX960 routers)* subsection under the *New Features in Junos OS Release for M Series, MX Series, and T Series Routers* main section of the Junos OS 11.4 Release Notes and Junos OS 12.1 Release Notes erroneously describes that the **show chassis fabric unreachable-destinations** command has been introduced. The correct command that has been introduced is the **show chassis fabric destinations** command, which is available in Junos OS Release 11.4R2 and later, and Junos OS Release 12.1R1 and later.

[*Release Notes*]

- The following additional information regarding the working of unnumbered interfaces applies to the *Example: Configuring an Unnumbered Ethernet Interface* section in the *Configuring an Unnumbered Interface* topic:

The sample configuration that is described works correctly on M Series and T Series routers. For unnumbered interfaces on MX Series routers, you must additionally configure static routes on an unnumbered Ethernet interface by including the **qualified-next-hop** statement at the **[edit routing-options static route destination-prefix]** hierarchy level to specify the unnumbered Ethernet interface as the next-hop interface for a configured static route.

[*Services Interfaces, Flow-Tap*]

- The following enhancements and additions apply to the *Example: Configuring Multichassis Link Aggregation in an Active-Active Bridging Domain on MX Series Routers* topic:
  - The *Topology Diagram* section fails to mention that interface **ge-1/0/2** functions as the ICCP link between the two PE devices, interface **ge-1/1/1** is the ICL-PL link, and interface **ge-1/1/4** is the link that connects to the server or the MC-LAG client device.
  - As a best practice, we recommend that you configure the ICCP and ICL interfaces over aggregated Ethernet interfaces instead of other interfaces such as Gigabit Ethernet interfaces, depending on your topology requirements and framework.
  - You must disable RSTP on the ICL-PL interfaces for an MC-LAG in an active-active bridging domain.
  - The *Step-by-Step Procedure* section for Router PE2 that is illustrated in the example is missing, although the quick configuration statements are presented.

To configure Router PE2:

1. Specify the number of aggregated Ethernet interfaces to be created.

```
[edit chassis]
user@PE2# set aggregated-devices ethernet device-count 5
```

2. Specify the members to be included within the aggregated Ethernet bundles.

```
[edit interfaces]
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae1
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae0
```

3. Configure the interfaces that connect to senders or receivers, the ICL interfaces, and the ICCP interfaces.

```
[edit interfaces]
user@PE2# set ge-1/0/3 flexible-vlan-tagging
user@PE2# set ge-1/0/3 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/3 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/3 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/4 flexible-vlan-tagging
user@PE2# set ge-1/0/4 encapsulation flexible-ethernet-services
user@PE2# set ge-1/0/4 unit 0 encapsulation vlan-bridge
user@PE2# set ge-1/0/4 unit 0 vlan-id-range 100-110
user@PE2# set ge-1/0/5 gigether-options 802.3ad ae0
user@PE2# set ge-1/1/0 gigether-options 802.3ad ae1
```

4. Configure parameters on the aggregated Ethernet bundles.

```
[edit interfaces ae0]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```

```
[edit interfaces ae1]
user@PE2# set flexible-vlan-tagging
user@PE2# set encapsulation flexible-ethernet-services
user@PE2# set unit 0 encapsulation vlan-bridge
user@PE2# set unit 0 vlan-id-range 100-110
user@PE2# set unit 0 multi-chassis-protection 100.100.100.1 interface ge-1/0/4.0
```

5. Configure LACP on the aggregated Ethernet bundles.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1
```

```
[edit interfaces ae1 aggregated-ether-options]
user@PE2# set lacp active
user@PE2# set lacp system-priority 100
user@PE2# set lacp system-id 00:00:00:00:00:05
user@PE2# set lacp admin-key 1
```

6. Configure the MC-LAG interfaces.

```
[edit interfaces ae0 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 5
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active
```

```
[edit interfaces ae1 aggregated-ether-options]
user@PE2# set mc-ae mc-ae-id 10
user@PE2# set mc-ae redundancy-group 10
user@PE2# set mc-ae chassis-id 1
user@PE2# set mc-ae mode active-active
user@PE2# set mc-ae status-control active
```

The multichassis aggregated Ethernet identification number (**mc-ae-id**) specifies which link aggregation group the aggregated Ethernet interface belongs to. The **ae0** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 5**. The **ae1** interfaces on Router PE1 and Router PE2 are configured with **mc-ae-id 10**.

The **redundancy-group 10** statement is used by ICCP to associate multiple chassis that perform similar redundancy functions and to establish a communication channel so that applications on peering chassis can send messages to each other. The **ae0** and **ae1** interfaces on Router PE1 and Router PE2 are configured with the same redundancy group **redundancy-group 10**.

The **chassis-id** statement is used by LACP for calculating the port number of the MC-LAG's physical member links. Router PE2 uses **chassis-id 1** to identify both its **ae0** and **ae1** interfaces. Router PE1 uses **chassis-id 0** to identify both its **ae0** and **ae1** interfaces.

The **mode** statement indicates whether an MC-LAG is in active-standby mode or active-active mode. Chassis that are in the same group must be in the same mode.

7. Configure a domain that includes the set of logical ports.

```
[edit bridge-domains bd0]
user@PE2# set domain-type bridge
user@PE2# set vlan-id all
user@PE2# set service-id 20
user@PE2# set interface ae0.0
user@PE2# set interface ae1.0
user@PE2# set interface ge-1/0/3.0
user@PE2# set interface ge-1/1/1.0
user@PE2# set interface ge-1/1/4.0
```

The ports within a bridge domain share the same flooding or broadcast characteristics in order to perform Layer 2 bridging.

The bridge-level **service-id** statement is required to link related bridge domains across peers (in this case Router PE1 and Router PE2), and should be configured with the same value.

8. Configure ICCP parameters.

```
[edit protocols iccp]
user@PE2# set local-ip-addr 100.100.100.2
user@PE2# set peer 100.100.100.1 redundancy-group-id-list 10
user@PE2# set peer 100.100.100.1 liveness-detection minimum-interval 1000
```

9. Configure the service ID at the global level.

```
[edit switch-options]
user@PE2# set service-id 10
```

You must configure the same unique network-wide configuration for a service in the set of PE routers providing the service. This service ID is required if the multichassis aggregated Ethernet interfaces are part of a bridge domain.

*[Network Interfaces, Ethernet Interfaces]*

- The following additional information regarding the compatibility of modules for the interoperation of RPM clients and RPM servers applies to the *Configuring RPM Probes* section in the *Configuring Real-Time Performance Monitoring* topic:

Keep the following points in mind when you configure RPM clients and RPM servers:

- You cannot configure an RPM client that is PIC-based and an RPM server that is based on either the Packet Forwarding Engine or Routing Engine to receive the RPM probes.
- You cannot configure an RPM client that is Packet Forwarding Engine-based and an RPM server that receives the RPM probes to be on the PIC or Routing Engine.
- The RPM client and RPM server must be located on the same type of module. For example, if the RPM client is PIC-based, the RPM server must also be PIC-based, and if the RPM server is Packet Forwarding Engine-based, the RPM client must also be Packet Forwarding Engine-based.

*[System Basics, Chassis-Level Features]*

- The *open-timeout* configuration statement topic and the *Configuring Default Timeout Settings for Services Interfaces* topic incorrectly state that the default value of the timeout period for TCP session establishment is 30 seconds. The correct default value is 5 seconds.

*[System Basics, Chassis-Level Features]*

- The Supported Platforms section of the *set chassis display message* command topic erroneously states that this command is supported on MX Series routers. This command is not available on MX Series routers.

*[System Basics, Chassis-Level Features]*

### **Infrastructure**

- The following additional information regarding the behavior of the *accept-data* statement for MC-LAG in an active-active bridge domain applies to the *Active-Active Bridging and VRRP over IRB Functionality on MX Series Routers Overview* topic:

For a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain and with VRRP configured over an integrated routing and bridging (IRB) interface, you must include the **accept-data** statement at the **[edit interfaces interface-name unit logical-unit-number family inet address address vrrp-group group-id]** hierarchy level to enable the router that functions as the master router to accept all packets destined for the virtual IP address.

On an MC-LAG, if you modify the source MAC address to be the virtual MAC address, you must specify the virtual IP address as the source IP address instead of the physical IP address. In such a case, the **accept-data** option is required for VRRP to prevent ARP from performing an incorrect mapping between IP and MAC addresses for customer edge (CE) devices. The **accept-data** attribute is needed for VRRP over IRB interfaces in MC-LAG to enable OSPF or other Layer 3 protocols and applications to work properly over multichassis aggregated Ethernet (mc-aeX) interfaces.

*[Network Interfaces, Ethernet Interfaces]*

- The following additional information regarding the configuration of peer IP addresses for ICCP peers and multichassis protection for MC-LAG applies to the *Configuring ICCP for MC-LAG* topic:

For Inter-Chassis Control Protocol (ICCP) in a multichassis link aggregation group (MC-LAG) configured in an active-active bridge domain, you must ensure that you configure the same peer IP address hosting the MC-LAG by including the **peer ip-address** statement at the **[edit protocols iccp]** hierarchy level and the **multi-chassis-protection peer ip-address** statement at the **[edit interfaces interface-name]** hierarchy level. Multichassis protection reduces the configuration at the logical interface level for MX Series routers with multichassis aggregated Ethernet (MC-AE) interfaces. If the ICCP is UP and the interchassis data link (ICL) comes UP, the router configured as standby will bring up the MC-AE interfaces shared with the peer active-active node specified by the **peer** statement.

For example, the following statements illustrate how the same peer IP address can be configured for both the ICCP peer and multichassis protection link:

```
set interfaces ae1 unit 0 multi-chassis-protection 10.255.34.112 interface ae0.0
set protocols iccp peer 10.255.34.112 redundancy-group-id-list 1
```

Although you can commit an MC-LAG configuration with various parameters defined for it, you can configure multichassis protection between two peers without configuring the ICCP peer address. You can also configure multiple ICCP peers and commit such a configuration.

*[Network Interfaces, Ethernet Interfaces]*

### ***Junos OS XML API and Scripting***

- The *NETCONF XML Management Protocol Guide* incorrectly states that when performing a confirmed commit operation using the `<commit>` element, the `<confirm-timeout>` value specifies the number of minutes for the rollback deadline. The value of the `<confirm-timeout>` element actually specifies the number of seconds for the rollback deadline.

[*NETCONF XML Management Protocol Guide*]

### ***J-Web Interface***

- To access the J-Web interface, your management device requires the following software:
  - Supported browsers—Microsoft Internet Explorer version 7.0 or Mozilla Firefox version 3.0
  - Language support—English-version browsers
  - Supported OS—Microsoft Windows XP Service Pack 3

### ***Layer 2 Ethernet Services***

- In the *Layer 2 Configuration Guide*, the examples provided in the sections, “Configuring Layer 2 Protocol Tunneling”, “Configuring BPDU Protection on Individual Interfaces”, and “Configuring BPDU Protection on All Edge Ports” are incorrect for configuring Layer 2 tunneling with routing instances.
- The following information regarding the differences in the default limit on MAC addresses that can be learned on an access port and a trunk port is inadvertently omitted from the *Limiting MAC Addresses Learned from an Interface in a Bridge Domain* topic:
  - For an access port, the default limit on the maximum number of MAC addresses that can be learned on an access port is 1024. Because an access port can be configured in only one bridge domain in a network topology, the default limit is 1024 addresses, which is same as the limit for MAC addresses learned on a logical interface in a bridge domain (configured by including the `interface-mac-limit limit` statement at the `[edit bridge-domains bridge-domain-name bridge-options interface interface-name]` or `[edit bridge-domains bridge-domain-name bridge-options]` hierarchy level.
  - For a trunk port, the default limit on the maximum number of MAC addresses that can be learned on a trunk port is 8192. Because a trunk port can be associated with multiple bridge domains, the default limit is the same as the limit for MAC addresses learned on a logical interface in a virtual switch instance (configured by including the `interface-mac-limit limit` statement at the `[edit routing-instances routing-instance-name switch-options interface interface-name]` for a virtual switch instance).

[*Layer 2 Configuration Guide, Bridging, Address Learning, and Forwarding*] ]

### **Multicast**

- The listings for the following RFCs incorrectly state that Junos OS supports only SSM include mode. Both include mode and exclude mode are supported in Junos OS Release 9.3 and later.
  - RFC 3376, *Internet Group Management Protocol, Version 3*
  - RFC 3590, *Source Address Selection for the Multicast Listener Discovery (MLD) Protocol* [Hierarchy and Standards Reference]

### **MPLS**

- **New output field in the show route command**—When the **no-propagate-ttl** statement is configured at the [edit protocols mpls] hierarchy level, a new line **NoPropagateTTL** is displayed in the output for the mpls.0 route table when you run the **show route** command.

### **Network Management**

- The *Supported Network Management Standards* topic fails to mention the following additional information:

On MX Series routers with MPC/MIC interfaces that use the ATM MIC with SFP, Junos OS substantially supports the following RFCs:

- RFC 5603, *PWE3 MIB*
- RFC 5601, *PW-FRAME-MIB*

[Junos OS Supported Standards]

- The documentation fails to clearly describe the characters that can be used for SNMPv3 authentication passwords. Besides numbers, uppercase letters, and lowercase letters, the following special characters are supported:

,./\<>;:'[]{}~!@#\$%^\*\_+=-`

In addition, the following special characters are also supported, but you must enclose them within quotation marks ("" ) if you enter them on the CLI; if you use a Network Management System to enter the password, the quotation marks are not required:

| & ( ) ?

The documentation also fails to clearly state that characters entered by simultaneously pressing the Ctrl key and additional keys are not supported. [PR/883083: This issue has been resolved]

- The syntax of the **filter-interfaces** statement in the *SNMP Configuration Statement* section is incorrect. The correct syntax is as follows:

```
filter-interfaces {
  all-internal-interfaces;
  interfaces interface-names{
    interface 1;
    interface 2;
```



```
}  
}
```

### ***Routing Policy and Firewall Filters***

- The following additional information regarding port mirroring functionality for IP-GRE tunneled traffic applies to the *Configuring Port Mirroring* topic

In the MPCs on M Series and MX Series routers, GRE and MPLS header information is not contained in the port-mirrored traffic corresponding to MPLS packets transmitted through IP-GRE tunnels.

[*Routing Policy*]

### ***Routing Protocols***

- In routing instances, when a BGP neighbor sends BGP messages to the local routing device, the incoming interface on which these messages are received must be configured in the same routing instance that the BGP neighbor configuration exists in. This is true for neighbors that are a single hop away or multiple hops away. [*Routing Protocols*]
- The following additional information regarding the behavior of MAC addresses in a VPLS dual-homed network with MSTP applies to the *Bridge Priority for Election of Root Bridge and Designated Bridge* topic:

Consider a sample scenario in which a dual-homed customer edge (CE) router is connected to two other provider edge (PE) routers, which function as the VPLS PE routers, with MTSP enabled on all these routers, and with the CE router operating as the root bridge. Integrated Routing and Bridging (IRB) interface is configured for the VPLS routing instances on the routers. In such a network, the MAC addresses that are learned in the VPLS domain continuously move between the LSI or virtual tunnel (VT) interfaces and the VPLS interfaces on both the PE routers. To avoid the continuous movement of the MAC addresses, you must configure root protection by including the **no-root-port** statement at the **[edit routing-instances routing-instance-name protocols mstp interface interface-name]** hierarchy level and configure the bridge priority as zero by including the **bridge priority 0** statement at the **[edit routing-instances routing-instance-name protocols mstp]** hierarchy level on the PE routers. This configuration on the PE routers is required to prevent the CE-side facing interfaces from becoming the route bridge.

[*Layer 2 Configuration Guide*]

- The *Supported MPLS Standards* topic fails to mention the following additional information:

On MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Junos OS substantially supports RFC 4385, *Pseudowire Emulation Edge-to-Edge (PWE3) Control Word for Use over an MPLS PSN*.

[*Junos OS Supported Standards*]

- The *Supported Carrier-of-Carriers and Interprovider VPN Standards* topic fails to mention the following additional information:

On MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Junos OS substantially supports the following RFCs:

- RFC 3985, *Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture*
- RFC 3916, *Requirements for Pseudo-Wire Emulation Edge-to-Edge (PWE3)*

[*Junos OS Supported Standards*]

- The *Supported IPv4, TCP, and UDP Standards* topic fails to mention the following additional information:

Junos OS substantially supports RFC 950, *Internet Standard Subnetting Procedure*.

[*Junos OS Supported Standards*]

- The *OSPF Configuration Guide* incorrectly includes the **transmit-interval** statement at the **[edit protocols ospf area area interface interface-name]** hierarchy level. The **transmit-interval** statement at this hierarchy level is deprecated in the Junos OS command-line interface.

[*OSPF Configuration Guide*]

### **Services Applications**

- The **rate** statement for packet sampling is now configured at the **[edit forwarding options sampling input family family]** hierarchy level.

[*Services Interfaces*]

- IPFIX sampling documentation did not reference the correct flow template. The documentation for “Configuring Inline Sampling” and “Configuring Inline Sampling for MX80 Routers” referred to the topic “Configuring Flow Aggregation to Use Version 9 Flow Templates” for information about sampling output, leading customers to believe that the IPv4 BGP\_NEXT\_HOP was supported for inline sampling. Inline sampling does not use Version 9 templates; they are used only for sampling done on a services PIC.

To view the correct flow template topic, “Configuring Flow Aggregation to Use IPFIX Flow Templates”, see [PR788037](#).

- The **aes-128-cbc**, **aes-192-cbc**, and **aes-256-cb** options that you can configure with the encryption-algorithm statement at the **[edit services ipsec-vpn ipsec proposal proposal-name]** hierarchy level are incorrectly specified as **ase-128-cbc**, **ase-192-cbc**, and **ase-256-cb** options in the following topics in the *Security Services* section of the *System Basics Configuration Guide*:

- *Security Services Configuration Statements*
- *Configuring Minimum IKE Requirements for IPsec on an ES PIC*
- *Configuring an IKE Proposal for Dynamic SAs*
- *encryption-algorithm*

[*System Basics, Security Services*]

- “The show services stateful-firewall flow-analysis command should be included in the System Basics and Services Command Reference Guide. This command displays stateful firewall flow statistics.”
- “The show services stateful-firewall subscriber-analysis command should be included in the System Basics and Services Command Reference Guide. This command displays information about the number of active subscribers on the service physical interface card (PIC).”
- In the *Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions Guide*, the section “Configuring Address Pools for Network Address Port Translation” should be revised as follows: The following variables should be added  
 Nr\_Addr\_PR\_Prefix – Number of usable pre-NAT IPv4 subscriber addresses in a “from” clause match condition  
 Nr\_Addr\_PU\_Prefix – Number of usable post-NAT IPv4 addresses configured in the NAT pool  
 Rounded\_Port\_Range\_Per\_IP =  $\text{ceil}[(\text{Nr\_Addr\_PR\_Prefix}/\text{Nr\_Addr\_PU\_Prefix})]$   
 \* Block\_Size The Forward Translation formulas should be: 1.  $\text{Pr\_Offset} = \text{Pr\_Prefix} - \text{Base\_Pr\_Prefix}$  2.  $\text{Pr\_Port\_Offset} = \text{Pr\_Offset} * \text{Block\_Size}$  3.  $\text{Rounded\_Port\_Range\_Per\_IP} = \text{ceil}[(\text{Nr\_Addr\_PR\_Prefix}/\text{Nr\_Addr\_PU\_Prefix})] * \text{Block\_Size}$  4.  $\text{Pu\_Prefix} = \text{Base\_Public\_Prefix} + \text{floor}(\text{Pr\_Port\_Offset}/\text{Rounded\_Port\_Range\_Per\_IP})$  5.  $\text{Pu\_Start\_Port} = \text{Pu\_Port\_Range\_Start} + (\text{Pr\_Port\_Offset} \% \text{Rounded\_Port\_Range\_Per\_IP})$  The Reverse Translation formulas should be: 1.  $\text{Pu\_Offset} = \text{Pu\_Prefix} - \text{Base\_Pu\_Prefix}$  2.  $\text{Pu\_Port\_Offset} = (\text{Pu\_Offset} * \text{Rounded\_Port\_Range\_Per\_IP}) + (\text{Pu\_Actual\_Port} - \text{Pu\_Port\_Range\_Start})$  3.  $\text{Subscriber\_IP} = \text{Base\_Pr\_Prefix} + \text{floor}(\text{Pu\_Port\_Offset} / \text{Block\_Size})$
- The following information should be added to the syntax of the **service-set (Services)** configuration statement topic in the *Services Interfaces Configuration Guide*. This information should appear under the **service-set service-set-name** level:

```

service-set-options {
    bypass-traffic-on-exceeding-flow-limits;
    bypass-traffic-on-pic-failure>;
    enable-asymmetric-traffic-processing;
    support-uni-directional-traffic;
}

```

This issue was being tracked by PR888803.

- The following information should be added after the second paragraph of the “Configuring Inline Sampling” topic in the *Services Interfaces Configuration Guide*:  
 The following limitations exist for inline sampling:
  - Flow records and templates cannot be exported if the flow collector is reachable through any management interface.
  - The flow collector should be reachable through the default routing table (inet.0 or inet6.0). If the flow collector is reachable via a non-default VPN routing and forwarding table (VRF), flow records and templates cannot be exported.
  - If the destination of the sampled flow is reachable through multiple paths, the IP\_NEXT\_HOP (Element ID 15) and OUTPUT\_SNMP (Element ID 14) in the IPv4 flow record would be set to the Gateway Address and SNMP Index of the first path seen in the forwarding table.

- If the destination of the sampled flow is reachable through multiple paths, the IP\_NEXT\_HOP (Element ID 15) and OUTPUT\_SNMP (Element ID 14) in the IPv6 flow records would be set to 0.
- The user-defined sampling instance has precedence over the global instance. When a user-defined sampling instance is attached to the FPC, the global instance is removed from the FPC, and the user-defined sampling instance is applied to the FPC.
- The Incoming Interface (IIF) and Outgoing Interface (OIF) should be part of the same VRF. If OIF is in a different VRF, DST\_MASK (Element ID 13), DST\_AS (Element ID 17), IP\_NEXT\_HOP (Element ID 15), and OUTPUT\_SNMP (Element ID 14) would be set to 0 in the flow records.
- Each Lookup Chip (LU) maintains and exports flows independent of other LUs. Traffic received on a media interface is distributed across all LUs in a multi-LU platform. It is likely that a single flow will be processed by multiple LUs. Therefore, each LU creates a unique flow and exports it to the flow collector. This can cause duplicate flows records to be seen on the flow collector. The flow collector should aggregate PKTS\_COUNT and BYTES\_COUNT for duplicate flow records to derive a single flow record.

This issue is being tracked by PR907991.

- The *System Basics and Services Command Reference* should include the following commands in the chapter “Dynamic Application Awareness Operational Mode Commands”:

**request services application-identification application:** Copy, disable, or enable a predefined application signature.

**request services application-identification group:** Copy, disable, or enable a predefined application signature group.

**show services application-identification application:** Display detailed information about a specified application signature, all application signatures, or a summary of the existing application signatures and nested application signatures. Both custom and predefined application signatures and nested application signatures can be displayed.

**show services application-identification group:** Display detailed or summary information about a specified application signature group or all application signature groups. Both custom and predefined application signature groups can be displayed.

**show services application-identification version:** Display the Junos OS application package version.

- The following command should appear in the network address operational mode commands:

```
clear services nat statistics
<interface interface-name>
  <service-set service-set-name>
```

The **<interface *interface-name*>** option clears NAT statistics for the specified interface only.

The `<service-set service-set-name>` option clears NAT statistics for the specified service set only.

The `clear services inline nat statistics` command should include the following option:

`<interface interface-name>`

The `<interface interface-name>` option clears inline NAT statistics for the specified interface only.

- The following additional information regarding the interoperation of sample actions in firewall filters and traffic sampling applies to the *Minimum Configuration for Traffic Sampling* section in the *Configuring Traffic Sampling* topic:

The following prerequisites apply to M, MX, and T Series routers when you configure traffic sampling on interfaces and in firewall filters:

- If you configure a sample action in a firewall filter for an inet or inet6 family on an interface without configuring the forwarding-options settings, operational problems might occur if you also configure port mirroring or flow-tap functionalities. In such a scenario, all the packets that match the firewall filter are incorrectly sent to the service PIC.
- If you include the `then sample` statement at the `[edit firewall family inet filter filter-name term term-name]` hierarchy level to specify a sample action in a firewall filter for IPv4 packets, you must also include the `family inet` statement at the `[edit forwarding-options sampling]` hierarchy level or the `instance instance-name family inet` statement at the `[edit forwarding-options sampling]` hierarchy level. Similarly, if you include the `then sample` statement at the `[edit firewall family inet6 filter filter-name term term-name]` hierarchy level to specify a sample action in a firewall filter for IPv6 packets, you must also include `family inet6` statement at the `[edit forwarding-options sampling]` hierarchy level or the `instance instance-name family inet6` statement at the `[edit forwarding-options sampling]` hierarchy level. Otherwise, a commit error occurs when you attempt to commit the configuration.
- Also, if you configure traffic sampling on a logical interface by including the sampling input or sampling output statements at the `[edit interface interface-name unit logical-unit-number]` hierarchy level, you must also include the `family inet | inet6` statement at the `[edit forwarding-options sampling]` hierarchy level, or the `instance instance-name family inet | inet6` statement at the `[edit forwarding-options sampling]` hierarchy level.

[*Services Interfaces, Traffic Sampling*]

- The *Configuring Port Mirroring* topic erroneously states that the `input` statement can be included under the `[edit forwarding-options port-mirroring family (inet | inet6) output]` hierarchy level. Only the `output` statement is available at the `[edit forwarding-options port-mirroring family (inet | inet6)]` hierarchy level. To configure the input packet properties for port mirroring, you must include the `input` statement at the `[edit forwarding-options port-mirroring]` hierarchy level.

To configure port mirroring on a logical interface, configure the following statements at the `[edit forwarding-options port-mirroring]` hierarchy level:

`[edit forwarding-options port-mirroring]`

```

input {
  maximum-packet-length bytes
  rate rate;
  run-length number;
}
family (inet|inet6) {
  output {
    interface interface-name {
      next-hop address;
    }
    no-filter-check;
  }
}

```

Also, the note incorrectly states that the **input** statement can also be configured at the **[edit forwarding-options port-mirroring]** hierarchy level and that it is only maintained for backward compatibility. The note also mentions that the configuration of the **output** statement is deprecated at the **[edit forwarding-options port-mirroring]** hierarchy level.

The correct behavior regarding the port-mirroring configuration for the packets to be mirrored and for the destination at which the packets are to be received is as follows:



**NOTE:** The **input** statement is deprecated at the **[edit forwarding-options port-mirroring family (inet | inet6)]** hierarchy level and is maintained only for backward compatibility. You must include the **input** statement at the **[edit forwarding-options port-mirroring]** hierarchy level.

[Services Interfaces, Port Mirroring]

- In the *Output Fields* section of the **show services ipsec-vpn ipsec security-associations** command topic, the descriptions of the **Local Identity** and **Remote Identity** fields are not clear and complete. The following are the revised descriptions of these fields:
  - **Local Identity**—Protocol, address or prefix, and port number of the local entity of the IPsec association. The format is **id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)**. The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the **id-data-len** parameter can be one of the following, depending on the address configured in the IPsec rule:
    - For an IPv4 address, the length is 4 and the value displayed is 3.
    - For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7.
    - For a range of IPv4 addresses, the length is 8 and the value displayed is 7.
    - For an IPv6 address prefix, the length is 16 and the value displayed is 15.
    - For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31.
    - For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31.

The value of the **id-data-presentation** field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the local peer of the IPsec association, it is displayed instead of the address details.

- **Remote Identity**—Protocol, address or prefix, and port number of the remote entity of the IPsec association. The format is **id-type-name (proto-name:port-number,[0..id-data-len] = iddata-presentation)**. The protocol is always displayed as any because it is not user-configurable in the IPsec rule. Similarly, the port number field in the output is always displayed as 0 because it is not user-configurable in the IPsec rule. The value of the **id-data-len** parameter can be one of the following, depending on the address configured in the IPsec rule:
  - For an IPv4 address, the length is 4 and the value displayed is 3.
  - For a subnet mask of an IPv4 address, the length is 8 and the value displayed is 7.
  - For a range of IPv4 addresses, the length is 8 and the value displayed is 7.
  - For an IPv6 address prefix, the length is 16 and the value displayed is 15.
  - For a subnet mask of an IPv6 address prefix, the length is 32 and the value displayed is 31.
  - For a range of IPv6 address prefixes, the length is 32 and the value displayed is 31.

The value of the **id-data-presentation** field denotes the IPv4 address or IPv6 prefix details. If the fully qualified domain name (FQDN) is specified instead of the address for the remote peer of the IPsec association, it is displayed instead of the address details.

[*Services Interfaces, IPsec Properties, Junos VPN Site Secure*]

- The *Supported Flow Monitoring and Discard Accounting Standards* topic fails to mention the following additional information:

On MX Series routers, Junos OS partially supports the following RFCs:

- RFC 5101, *Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of IP Traffic Flow Information*
- RFC 5102, *Information Model for IP Flow Information Export*

[*Junos OS Supported Standards*]

- The following additional information applies to the sample configuration described in the *Example: Flow-Tap Configuration* topic of the *Flow Monitoring* chapter.



.....

**NOTE:** The described example applies only to M Series and T Series routers, except M160 and TX Matrix routers. For MX Series routers, because the flow-tap application resides in the Packet Forwarding Engine rather than a service PIC or Dense Port Concentrator (DPC), the Packet Forwarding Engine must send the packet to a tunnel logical (vt-) interface to encapsulate the intercepted packet. In such a scenario, you need to allocate a tunnel interface and assign it to the dynamic flow capture process for FlowTapLite to use.

.....

[*Services Interfaces, Flow Monitoring*]

- The topic *Configuring Secured Port Block Allocation* contains a note listing configuration changes that require a reboot of the services PIC. The note has been updated to include change to the NAT pool name.

### **Subscriber Access Management**

- The **show subscribers** topic in the *Junos OS System Basics and Services Command Reference* omits the following information about using the **address** option for the **show subscribers** command.

When you issue the **show subscribers address** command, you must specify the IPv4 or IPv6 address prefix *without* a netmask, as shown in the following example:

```
user@host> show subscribers address 192.168.17.1 detail
```

If you specify the IP address as a prefix *with* a netmask, as shown in the following example, the router displays a message that the IP address is invalid, and rejects the command:

```
user@host> show subscribers address 192.168.17.1/32 detail
Invalid argument: invalid ip_address 192.168.17.1/32
```

[*Junos OS System Basics and Services Command Reference*]

- The “Configuring Per-Subscriber Session Accounting” topic in the *Subscriber Access Configuration Guide* incorrectly states that the **update-interval** statement rounds up an interval of 10 through 15 minutes to 15. The actual behavior is that all configured values are rounded up to the next higher multiple of 10. For example, the values 811 through 819 are all accepted by the CLI, but are all rounded up to 820.

[*Subscriber Access*]

- The “DHCP in Broadband Networks” topic erroneously states that the Junos OS subscriber management solution currently supports only DHCP as a multiple-client configuration protocol. However, subscriber management solutions support DHCP and PPPoE as multiple-client configuration protocols.

[*Broadband Subscriber Management Solutions*]

- The “Configuring Service Packet Counting” topic in the *Junos OS Subscriber Access Configuration Guide* does not include the following configuration guideline. When you specify the **service-accounting** action for the term, you cannot additionally configure the **count** action in the same term.



*[Subscriber Access]*

- The table titled “Supported Juniper Networks VSAs” in the “Juniper Networks VSAs Supported by the AAA Service Framework” topic lists RADIUS VSA 26-157 (IPv6-NdRa-Pool-Name). This VSA is not supported and should not appear in the table.

*[Subscriber Access]*

- The “Configuring a Dynamic Profile for Client Access” topic erroneously uses the **\$junos-underlying-interface** variable when an IGMP interface is configured 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
```

- 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.

*[Subscriber Access, System Basics]*

- When you modify dynamic CoS parameters with a RADIUS change of authorization (CoA) message, Junos OS accepts invalid configurations. For example, if you specify a transmit rate that exceeds the allowed 100 percent, the system does not reject the configuration and returns unexpected shaping behavior.

*[Subscriber Access]*

- Juniper Networks does not support multicast RIF mapping and ANCP when configured simultaneously on the same logical interface. For example, configuring a multicast VLAN and ANCP on the same logical interface is not supported, and the subscriber VLANs are the same for both ANCP and multicast.

*[Subscriber Access]*

- The *Subscriber Access Configuration Guide* incorrectly describes the **authentication-order** statement as it is used for subscriber access management. When configuring the **authentication-order** statement for subscriber access management, you must always specify the **radius** method. Subscriber access management does not support the **password** keyword (the default), and authentication fails when you do not specify an authentication method.

*[Subscriber Access]*

- In the *Subscriber Access Configuration Guide*, the “Juniper Networks VSAs Supported by the AAA Service Framework” table and the “RADIUS-Based Mirroring Attributes” table incorrectly describe VSA 26-59. The correct description is as follows:

Attribute Number	Attribute Name	Description
26-59	Med-Dev-Handle	Identifier that associates mirrored traffic to a specific subscriber.

*[Subscriber Access]*

- In the *Subscriber Access Configuration Guide*, the table titled "Supported Juniper Networks VSAs" in the "Juniper Networks VSAs Supported by the AAA Service Framework" topic lists RADIUS VSA 26-42 (Input-Gigapackets) and VSA 26-43 (Output-Gigapackets). These two VSAs are not supported.

[*Subscriber Access*]

- In the *Junos OS Subscriber Access Configuration Guide*, the "Qualifications for Change of Authorization" section in the topic titled "RADIUS-initiated Change of Authorization (CoA) Overview", has been rewritten as follows to clarify how CoA uses the RADIUS attributes and VSAs.

### Qualifications for Change of Authorization

To complete the change of authorization for a user, you specify identification attributes and session attributes. The identification attributes identify the subscriber. Session attributes specify the operation (activation or deactivation) to perform on the subscriber's session and also include any client attributes for the session (for example, QoS attributes). The AAA Service Framework handles the actual request.

Table 7 on page 315 shows the identification attributes for CoA operations.



**NOTE:** Using the Acct-Session-ID attribute to identify the subscriber session is more explicit than using the User-Name attribute. When you use the Acct-Session-ID, the attribute identifies the specific subscriber and session. When you use the User-Name as the identifier, the CoA operation is applied to the first session that was logged in with the specified username. However, because a subscriber might have multiple sessions associated with the same username, the first session might not be the correct session for the CoA operation.

**Table 7: Identification Attributes**

Attribute	Description
User-Name [RADIUS attribute 1]	Subscriber username.
Acct-Session-ID [RADIUS attribute 44]	Specific subscriber and session.

Table 8 on page 315 shows the session attributes for CoA operations. Any additional client attributes that you include depend on your particular session requirements.

**Table 8: Session Attributes**

Attribute	Description
Activate-Service [Juniper Networks VSA 26–65]	Service to activate for the subscriber.
Deactivate-Service [Juniper Networks VSA 26–66]	Service to deactivate for the subscriber.

[Subscriber Access]

- The “overhead-accounting” configuration statement topic should include the following note under the **cell-mode** option:



**NOTE:** Cell mode is supported only on logical interfaces and interface sets; it is not supported on physical interfaces (ifd or ifd-remaining).

[Subscriber Access Configuration Guide, Class of Service Configuration Guide]

- The *Junos OS Subscriber Management Scaling Values (XLS)* spreadsheet erroneously states that the maximum number of PPPoE interfaces per MPC1 is 15,996. The correct value is 31,998.

[Subscriber Management Scaling]

- The documentation for the subscriber management domain mapping feature in the *Subscriber Access Configuration Guide* describes using the **aaa-logical-system** and **target-logical-system** statements to configure mapping to a non-default logical system. Subscriber management is supported in the default logical system only. Configuring a non-default logical system for subscriber management is not supported in current Junos OS releases.

[Subscriber Access]

- The *Example: HTTP Service Attached to a Static Interface* topic in the *Junos OS Subscriber Access Configuration Guide* provides an incorrect example for configuring a service filter as a walled garden. The correct example is as follows:

The following example uses a service filter as a walled garden by defining a rule named `redirect`, referencing the rule in a profile named `http-redirect`, configuring a service set named `http-redirect` that references the `http-redirect` captive portal content delivery profile, and attaching the `http-redirect` service set to static interface `ge-1/0/1.0`.

```
[edit services]
captive-portal-content-delivery {
  rule redirect {
    match-direction input;
    term t1 {
      from {
        destination-address {
          100.0.1.1/32;
        }
      }
      then {
        redirect http://www.google.com;
      }
    }
  }
  profile http-redirect {
    cpcd-rules redirect;
  }
}
service-set http-redirect {
  captive-portal-content-delivery-profile http-redirect;
  interface-service {
    service-interface ms-1/0/0;
  }
}

[edit interfaces ge-1/0/1]
unit 0 {
  family inet {
    service {
      input {
        service-set http-redirect service-filter walled;
```

```

    }
    output {
        service-set http-redirect;
    }
}
address 10.1.3.2/24;
}
}

```

[Subscriber Access]

- The *Dedicated Queue Scaling for CoS Configurations on Trio MPC/MIC Interfaces Overview* topic in the *Junos OS Subscriber Access Configuration Guide* does not explain the queuing behavior on 30-Gigabit Ethernet Queuing MPCs with only one MIC. See [Dedicated Queue Scaling for CoS Configurations on MIC and MPC Interfaces Overview](#) for a more complete explanation of dedicated queue scaling.
- In the *Subscriber Access Configuration Guide*, there is an error in the *Example: Configuring RADIUS-Based Subscriber Authentication and Accounting* topic. In the example, the **profile** stanza incorrectly includes the **authentication** statement. The correct statement is **authentication-order**, as shown in the following sample:

```

profile isp-bos-metro-fiber-basic {
    authentication-order radius;
}

```

[Subscriber Access]

- The *MX Series 3D Universal Edge Router Interface Module Reference* does not state that VLAN demux configurations are not supported on MX Series routers that have any of the following line cards installed:
  - Enhanced Queuing Ethernet Services DPCs (DPCE-X-Q)
  - Enhanced Queuing IP Services DPCs (DPCE-R-Q)

The nonsupport includes any configuration stacked on top of a VLAN demux. For example, although PPPoE is supported, PPPoE over aggregated Ethernet interfaces is not supported when one of these cards is installed, because this configuration requires PPPoE to be stacked on a VLAN demux.

- The Subscriber Management Scaling Values spreadsheet incorrectly reports the maximum number of L2TP LAC sessions per chassis entry for MX240, MX480, and MX960 routers for Junos OS Release 12.2. The maximum number of L2TP LAC sessions per chassis is 32,000 with the RE-S-2000 and 128,000 with the RE-S-1800. Previously this was reported as 60,000, with no distinction between Routing Engine models.
- In the *AAA Service Framework Feature Guide for Subscriber Management*, the **parse-direction** (Domain Map) statement and the *Specifying the Parsing Direction for Domain Names* topic show an incorrect default setting for the **parse-direction** statement. The correct default is the **left-to-right** direction.
- The *Example: HTTP Service Within a Service Set* topic in the *Subscriber Access Configuration Guide* erroneously describes how to configure captive portal content delivery rules in service sets.

Use the following procedure to configure captive portal content delivery rules in service sets:

1. Define one or more rules with the **rule rule-name** statement at the **[edit services captive-portal-content-delivery]** hierarchy level. In each rule you specify one or more terms to match on an application, destination address, or destination prefix list; where the match takes place; and actions to be taken when the match occurs,
  2. (Optional) Define one or more rule sets by listing the rules to be included in the set with the **rule-set rule-set-name** statement at the **[edit services captive-portal-content-delivery]** hierarchy level.
  3. Configure a captive portal content delivery profile with the **profile profile-name** statement at the **[edit services captive-portal-content-delivery]** hierarchy level.
  4. In the profile, specify a list of rules with the **cpcd-rules [rule-name]** statement or a list of rule sets with the **cpcd-rule-sets [rule-set-name]** statement. Both statements are at the **[edit services captive-portal-content-delivery profile profile-name]** hierarchy level.
  5. Associate the profile with a service set with the **captive-portal-content-delivery-profile profile-name** statement at the **[edit services service-set service-set-name]** hierarchy level.
- In the *Junos OS Subscriber Access Feature Guide*, the **fail-over-within-preference** statement at the **[edit services l2tp]** hierarchy level is incorrectly spelled. The correct spelling for this statement is **failover-within-preference**.
  - The *LAC Tunnel Selection Overview* topic in the *Subscriber Access Configuration Guide* incorrectly describes the current behavior for failover between preference levels. The topic states that when the tunnels at every preference level have a destination in the lockout state, the LAC cycles back to the highest preference level and waits for the lockout time for a destination at that level to expire before attempting to connect and starting the process over.

In fact, the current behavior in this situation is that from the tunnels present at the lowest level of preference (highest preference number), the LAC selects the tunnel that has the destination with the shortest remaining lockout time. The LAC ignores the lockout and attempts to connect to the destination.

- The *LAC Tunnel Selection Overview*, *Configuring Weighted Load Balancing for LAC Tunnel Sessions* and *weighted-load-balancing (L2TP LAC)* topics in the *Junos OS Subscriber Access Configuration Guide* incorrectly describe how weighted load balancing works on an L2TP LAC. The topics state that the tunnel with the highest weight (highest session limit) within a preference level is selected until it has reached its maximum sessions limit, and then the tunnel with the next higher weight is selected, and so on.

In fact, when weighted load balancing is configured, tunnels are selected randomly within a preference level, but the distribution of selected tunnels is related to their weight. The LAC generates a random number within a range equal to the aggregate total of all session limits for all tunnels in the preference level. Portions of the range—pools of numbers—are associated with the tunnels according to their weight;

a higher weight results in a larger pool. The random number is more likely to be in a larger pool, so a tunnel with a higher weight (larger pool) is more likely to be selected than a tunnel with a lower weight (smaller pool).

For example, consider a level that has only two tunnels, A and B. Tunnel A has a maximum sessions limit of 1000 and tunnel B has a limit of 2000 sessions, resulting in an aggregate total of 3000 sessions. The LAC generates a random number in the range from 0 through 2999. A pool of 1000 numbers, the portion of the range from 0 through 999, is associated with tunnel A. A pool of 2000 numbers, the portion of the range from 1000 through 2999, is associated with tunnel B. If the generated number is less than 1000, then tunnel A is selected, even though it has a lower weight than tunnel B. If the generated number is 1000 or larger, then tunnel B is selected. Because the pool of possible generated numbers for tunnel B (2000) is twice that for tunnel A (1000), tunnel B is, *on average*, selected twice as often as tunnel A.

- The table in topic, “AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS,” incorrectly indicates that VSA 26-1 (Virtual-Router) supports CoA Request messages. VSA 26-1 does not support CoA Request messages.

#### ***Timing and Synchronization***

- The *Supported Time Synchronization Standards* topic fails to mention the following additional information:

On MX Series routers with the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP, Junos OS substantially supports RFC 4553, *Structure-Agnostic Time Division Multiplexing (TDM) over Packet (SAToP)*

[*Junos OS Supported Standards*]

#### ***User Interface and Configuration***

- The **show system statistics bridge** command displays system statistics on MX Series routers.

[*System Basics Command Reference*]

- The note in the *Installing the J-Web Software* topic that mentions that M Series or T Series routers must be running Junos OS version 7.3 or later to support the J-Web interface is incomplete. The following note accurately describes the support of the J-Web interface on M Series and T series routers.

M Series routers or T320, T640, and TX Matrix routers must be running Junos OS version 7.3 or later to support the J-Web interface. Except the T320, T640, and TX Matrix routers, other T Series routers do not support the J-Web software.

[*J-Web Interface User Guide*]

### VPNs

- In “Chapter 19, Configuring VPLS” of the *VPNs Configuration Guide*, an incorrect statement that caused contradictory information about which platforms support LDP BGP interworking has been removed. The M7i router was also omitted from the list of supported platforms. The M7i router does support LDP BGP interworking.

[VPNs]

- The **l3vpn** statement documentation states that this statement is not supported on MX Series routers with both MS-DPCs and MPCs installed. However, it should state that the **l3vpn** statement is not supported on MX Series routers with both DPCs and MPCs installed.

[VPNs]

- The following guideline regarding the support of LSI traffic statistics on M Series routers is missing from the *General Limitations on IP-Based Filtering* section in the *Filtering Packets in Layer 3 VPNs Based on IP Headers* topic:

Label-switched interface (LSI) traffic statistics are not supported for Intelligent Queuing 2 (IQ2), Enhanced IQ (IQE), and Enhanced IQ2 (IQ2E) PICs on M Series routers.

[VPNs, Layer 3 VPNs]

- The following limitation regarding firewall filters configured in conjunction with the **vrf-table-label** statement is missing from the *General Limitations on IP-Based Filtering* in the *Filtering Packets in Layer 3 VPNs Based on IP Headers* topic:

Firewall filters cannot be applied to interfaces included in a routing instance on which you have configured the **vrf-table-label** statement.

This documentation is applicable to all J Series, M Series, T Series, and SRX Series routers.

[VPNs, Layer 3 VPNs]

- The descriptions of the **pw-label-ttl-1** and **router-alert-label** options in the *control-channel (Protocols OAM)* configuration statement topic are incorrectly and interchangeably stated. The correct descriptions of these options are as follows:
  - **pw-label-ttl-1**—For BGP-based pseudowires that send OAM packets with the MPLS pseudowire label and time-to-live (TTL) set to 1.
  - **router-alert-label**—For BGP-based pseudowires that send OAM packets with router alert label.

[VPNs, Layer 2 VPNs]



## Changes to the Junos OS Documentation Set

---

The following are the changes made to the Junos OS documentation set:

- Carrier-grade NAT and softwire documentation is no longer included in the *Junos OS Services Configuration Guide*. The documentation is now available at the following subject-based web page: Next-Generation Network Addressing Carrier-Grade NAT and IPv6 Solutions—[http://www.juniper.net/techpubs/en\\_US/junos12.1/information-products/pathway-pages/ngna-solutions/next-generation-network-addressing-solutions.html](http://www.juniper.net/techpubs/en_US/junos12.1/information-products/pathway-pages/ngna-solutions/next-generation-network-addressing-solutions.html)
- The documentation for ukernel and JSF supported Application Layer Gateways (ALGs) has been substantially re-written, and is available at the following web pages:
  - *ALG Descriptions* (ukernel)—[http://www.juniper.net/techpubs/en\\_US/junos12.1/topics/concept/alg-descriptions.html](http://www.juniper.net/techpubs/en_US/junos12.1/topics/concept/alg-descriptions.html)
  - *ALG Descriptions* (JSF)—[http://www.juniper.net/techpubs/en\\_US/junos12.1/topics/concept/alg-descriptions-jsf.html](http://www.juniper.net/techpubs/en_US/junos12.1/topics/concept/alg-descriptions-jsf.html)
- Stateless firewall filter and traffic policer documentation is no longer included in the *Junos OS Policy Framework Configuration Guide*. This material is now available in the *Routing Policy Configuration Guide* only.
- Routing policy, traffic sampling, forwarding, and monitoring documentation is no longer included in the *Junos OS Policy Framework Configuration Guide*. This material is now available in the *Junos OS Routing Policy Configuration Guide*.

In addition, individual HTML pages have a **Print** link in the upper left corner of the text area on the page.

- A new topic, *CGN Implementation: Best Practices*, which provides experience-based recommendations for configuring carrier-grade NAT, has been added to the documentation set. The new topic is available at [http://www.juniper.net/techpubs/en\\_US/junos12.2/topics/concept/nat-best-practices.html](http://www.juniper.net/techpubs/en_US/junos12.2/topics/concept/nat-best-practices.html)
- ALG documentation for MX Series platforms has been updated. The topic has been reorganized and expanded, with particular emphasis on SIP and SIP-NAT interaction. An updated version of the documentation is available at the following PR link location: [PR817816](#)

### Related Documentation

- [New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 100](#)
- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 154](#)
- [Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 170](#)
- [Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 171](#)
- [Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 322](#)

## Upgrade and Downgrade Instructions for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers

This section discusses the following topics:

- [Basic Procedure for Upgrading to Release 12.2 on page 322](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 325](#)
- [Upgrading a Router with Redundant Routing Engines on page 326](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 326](#)
- [Upgrading the Software for a Routing Matrix on page 328](#)
- [Upgrading Using Unified ISSU on page 329](#)
- [Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR on page 329](#)
- [Downgrading from Release 12.2 on page 330](#)

---

### Basic Procedure for Upgrading to Release 12.2

In order to upgrade to Junos OS 10.0 or later, you must be running Junos OS 9.0S2, 9.1S1, 9.2R4, 9.3R3, 9.4R3, 9.5R1, or later minor versions, or you must specify the **no-validate** option on the **request system software install** command.

When upgrading or downgrading Junos OS, always use the **jinstall** package. Use other packages (such as the **bundle** 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:** With Junos OS Release 9.0 and later, the compact flash disk memory requirement for Junos OS 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 OS 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 Junos OS. 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 OS System Basics Configuration Guide*.

.....

The download and installation process for Junos OS Release 12.2 is different from previous Junos OS releases. Follow these steps:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks Web page:  
<http://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. 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-12.2R9-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-12.2R9-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.

Including 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 OS Release 12.2 **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 about how to monitor call traffic before upgrading, see the *Junos OS Multiplay Solutions Guide*.

---

### Upgrade and Downgrade Support Policy for Junos OS Releases

---

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

### Upgrading a Router with Redundant Routing Engines

---

If the router has two Routing Engines, perform a Junos OS 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 OS 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 Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1

---

In releases prior to Junos OS Release 10.1, the draft-rosen multicast VPN feature implements the unicast **lo0.x** address configured within that instance as the source address used to establish PIM neighbors and create the multicast tunnel. In this mode, the multicast VPN loopback address is used for reverse path forwarding (RPF) route resolution to create the reverse path tree (RPT), or multicast tunnel. The multicast VPN loopback address is also used as the source address in outgoing PIM control messages.

In Junos OS Release 10.1 and later, you can use the router's main instance loopback (**lo0.0**) address (rather than the multicast VPN loopback address) to establish the PIM state for the multicast VPN. We strongly recommend that you perform the following procedure when upgrading to Junos OS Release 10.1 if your draft-rosen multicast VPN network includes both Juniper Network routers and other vendors' routers functioning as provider edge (PE) routers. Doing so preserves multicast VPN connectivity throughout the upgrade process.

Because Junos OS Release 10.1 supports using the router's main instance loopback (**lo0.0**) address, it is no longer necessary for the multicast VPN loopback address to match the main instance loopback address **lo0.0** to maintain interoperability.



**NOTE:** You might want to maintain a multicast VPN instance **lo0.x** address to use for protocol peering (such as IBGP sessions), or as a stable router identifier, or to support the PIM bootstrap server function within the VPN instance.

Complete the following steps when upgrading routers in your draft-rosen multicast VPN network to Junos OS Release 10.1 if you want to configure the routers's main instance loopback address for draft-rosen multicast VPN:

1. Upgrade all M7i and M10i routers to Junos OS Release 10.1 before you configure the loopback address for draft-rosen Multicast VPN.



**NOTE:** Do not configure the new feature until all the M7i and M10i routers in the network have been upgraded to Junos OS Release 10.1.

2. After you have upgraded all routers, configure each router's main instance loopback address as the source address for multicast interfaces.

Include the **default-vpn-source interface-name loopback-interface-name** statement at the **[edit protocols pim]** hierarchy level.

3. After you have configured the router's main loopback address on each PE router, delete the multicast VPN loopback address (**lo0.x**) from all routers.

We recommend that you remove the multicast VPN loopback address from all PE routers from other vendors. In Junos OS releases prior to 10.1, to ensure interoperability with other vendors' routers in a draft-rosen multicast VPN network, you had to perform additional configuration. Remove that configuration from both the Juniper Networks routers and the other vendors' routers. This configuration should be on Juniper Networks routers and on the other vendors' routers where you configured the **lo0.mvpn** address in each VRF instance as the same address as the main loopback (**lo0.0**) address.

This configuration is not required when you upgrade to Junos OS Release 10.1 and use the main loopback address as the source address for multicast interfaces.



**NOTE:** To maintain a loopback address for a specific instance, configure a loopback address value that does not match the main instance address (**lo0.0**).

For more information about configuring the draft-rosen Multicast VPN feature, see the *Junos OS Multicast Configuration 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 OS 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 OS CLI by using the **lcc** option). To avoid network disruption during the upgrade, ensure that the following conditions are met 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 in the routing matrix, use the CLI **show system storage** command. To determine the amount of DRAM currently available on all 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 Junos OS 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 be 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 OS 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 OS on the backup Routing Engine.
4. 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.
5. Install the new software on the new backup Routing Engine (re0).

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

---

### Upgrading Using Unified ISSU

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

---

### Upgrading from Junos OS Release 9.2 or Earlier on a Router Enabled for Both PIM and NSR

Junos OS 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 OS 9.3 Release 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 unified ISSU procedures described in the other sections of these instructions.

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

To disable and reenale PIM:

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

[edit]

```
user@host# deactivate protocols pim
user@host# commit
```

2. Upgrade to Junos OS 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 unified ISSU.

3. After the router reboots and is running the upgraded Junos OS, enter configuration mode, disable PIM NSR with the **nonstop-routing disable** statement, and then reenale PIM.

[edit]

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

---

### Downgrading from Release 12.2

To downgrade from Release 12.2 to another supported release, follow the procedure for upgrading, but replace the 12.2 **jinstall** package with one that corresponds to the appropriate release.



**NOTE:** You cannot downgrade more than three releases. For example, if your routing platform is running Junos OS Release 11.4, you can downgrade the software to Release 10.4 directly, but not to Release 10.3 or earlier. As a workaround, you can first downgrade to Release 10.4 and then downgrade to Release 10.3.

---

For more information, see the [Junos OS Installation and Upgrade Guide](#).

#### Related Documentation

- [New Features in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 100](#)

- [Changes in Default Behavior and Syntax in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 154](#)
- [Known Behavior in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 170](#)
- [Issues in Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 171](#)
- [Errata and Changes in Documentation for Junos OS Release 12.2 for M Series, MX Series, and T Series Routers on page 285](#)

## Junos OS Documentation and Release Notes

---

For a list of related Junos OS 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 *Junos OS 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](mailto: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 or topic name
- URL or page number
- Software release version (if applicable)

## 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

---

5 August 2015—Revision 4, Junos OS 12.2R9 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

6 May 2015—Revision 3, Junos OS 12.2R9 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

16 September 2014—Revision 2, Junos OS 12.2R9 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

9 September 2014—Revision 1, Junos OS 12.2R9 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

29 April 2014—Revision 3, Junos OS 12.2R8 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

22 April 2014—Revision 2, Junos OS 12.2R8 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

15 April 2014—Revision 1, Junos OS 12.2R8 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

13 February 2014—Revision 4, Junos OS 12.2R7 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

22 January 2014—Revision 3, Junos OS 12.2R7 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

16 January 2014—Revision 2, Junos OS 12.2R7 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

8 January 2014—Revision 1, Junos OS 12.2R7 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

19 November 2013—Revision 6, Junos OS 12.2R6 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

24 October 2013—Revision 5, Junos OS 12.2R6 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

21 October 2013—Revision 4, Junos OS 12.2R6 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

18 October 2013—Revision 3, Junos OS 12.2R6 – ACX Series, and the M Series, MX Series, and T Series.

11 October 2013—Revision 2, Junos OS 12.2R6 – ACX Series, and the M Series, MX Series, and T Series.

10 October 2013—Revision 1, Junos OS 12.2R6 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

17 September 2013—Revision 6, Junos OS 12.2R5 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

08 August 2013—Revision 5, Junos OS 12.2R5 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

31 July 2013—Revision 4, Junos OS 12.2R5 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

24 July 2013—Revision 3, Junos OS 12.2R5 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

17 July 2013—Revision 2, Junos OS 12.2R5 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

10 July 2013—Revision 1, Junos OS 12.2R5 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

20 June 2013—Revision 6, Junos OS 12.2R4 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

30 May 2013—Revision 5, Junos OS 12.2R4 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

07 May 2013—Revision 4, Junos OS 12.2R4 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

25 April 2013—Revision 3, Junos OS 12.2R4 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

18 April 2013—Revision 2, Junos OS 12.2R4 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

9 April 2013—Revision 1, Junos OS 12.2R4 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

21 February 2013—Revision 3, Junos OS 12.2R3 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

31 January 2013—Revision 2, Junos OS 12.2R3 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

29 January 2013—Revision 1, Junos OS 12.2R3 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

08 January 2013—Revision 2, Junos OS 12.2R2 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

20 November 2012—Revision 1, Junos OS 12.2R2 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

02 October 2012—Revision 3, Junos OS 12.2R1 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

24 September 2012—Revision 2, Junos OS 12.2R1 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

05 September 2012—Revision 1, Junos OS 12.2R1 – ACX Series, EX Series, and the M Series, MX Series, and T Series.

Copyright © 2015, 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.