

# Release Notes: Junos<sup>®</sup> OS Release 15.1F7 for the MX Series

15.1F7  
27 February 2018

## Contents

Introduction .....	4
Junos OS Release Notes for MX Series 3D Universal Edge Routers .....	4
New and Changed Features .....	5
Hardware .....	5
Class of Service (CoS) .....	10
Forwarding and Sampling .....	11
General Routing .....	11
High Availability and Resiliency .....	13
Interfaces and Chassis .....	13
IPv6 .....	17
Management .....	17
MPLS .....	17
Multicast .....	18
Network Management and Monitoring .....	18
Routing Protocols .....	19
Services Applications .....	22
Software Installation and Upgrade .....	24
Software-Defined Networking .....	24
Subscriber Management and Services .....	27
System Logging .....	34
VPNs .....	35
Changes in Behavior and Syntax .....	35
Authentication, Authorization and Accounting (AAA) (RADIUS) .....	37
Flow-based and Packet-based Processing .....	37
General Routing .....	37
High Availability (HA) and Resiliency .....	37
Interfaces and Chassis .....	37
IPv6 .....	38
MPLS .....	38
Network Management and Monitoring .....	38
Routing Policy and Firewall Filters .....	39

Routing Protocols	39
Services Applications	40
Subscriber Management and Services (MX Series)	41
System Logging	43
System Management	49
Platform and Infrastructure	49
Virtual Chassis	51
VPNs	51
Known Behavior	51
Hardware	52
General Routing	52
Interfaces and Chassis	54
MPLS	55
OpenFlow	55
Services Applications	55
Subscriber Management and Services (MX Series)	55
User Interface and Configuration	56
Known Issues	56
Class of Service (CoS)	57
Forwarding and Sampling	57
General Routing	58
High Availability (HA) and Resiliency	63
Infrastructure	63
Interfaces and Chassis	63
Layer 2 Features	64
Layer 2 Ethernet Services	64
Multiprotocol Label Switching (MPLS)	65
Platform and Infrastructure	66
Routing Protocols	67
Services Applications	69
Software Installation and Upgrade	69
Subscriber Access Management	69
User Interface and Configuration	70
VPNs	70
Resolved Issues	70
Resolved Issues: 15.1F7	71
Resolved Issues: 15.1F6	92
Resolved Issues: 15.1F5	127
Resolved Issues: 15.1F4	141
Resolved Issues: 15.1F3	147
Resolved Issues: 15.1F2	156
Documentation Updates	166
Subscriber Management Provisioning Guide	166
Migration, Upgrade, and Downgrade Instructions	166
Basic Procedure for Upgrading to Release 15.1F7	167
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)	169
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)	170
Upgrade and Downgrade Support Policy for Junos OS Releases	172

---

Upgrading a Router with Redundant Routing Engines . . . . .	173
Upgrading Using Unified ISSU . . . . .	173
Downgrading from Release 15.1 . . . . .	174
Product Compatibility . . . . .	174
Hardware Compatibility . . . . .	174
Third-Party Components . . . . .	176
Finding More Information . . . . .	176
Documentation Feedback . . . . .	176
Requesting Technical Support . . . . .	177
Self-Help Online Tools and Resources . . . . .	177
Opening a Case with JTAC . . . . .	177
Revision History . . . . .	178

## Introduction

---

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric systems, QFX Series, SRX Series, T Series, and Junos Fusion.

The Junos OS 15.1F7 release is an MX Series only release. These release notes accompany Junos OS Release 15.1F7 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

## Junos OS Release Notes for MX Series 3D Universal Edge Routers

---

These release notes accompany Junos OS Release 15.1F7 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at

[https://www.juniper.net/documentation/en\\_US/release-independent/junos/information-products/pathway-pages/junos/product/index.html](https://www.juniper.net/documentation/en_US/release-independent/junos/information-products/pathway-pages/junos/product/index.html)



**CAUTION:** This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for MX Series. We do not recommend using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 15.1.



**NOTE:** The Junos OS 15.1F7 release does not support the following MX Series routers or features:

- Virtual MX Series router (vMX)
- BBE support for the Broadband Network Gateway (BNG)

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 35](#)
- [Known Behavior on page 51](#)
- [Known Issues on page 56](#)
- [Resolved Issues on page 70](#)
- [Documentation Updates on page 166](#)
- [Migration, Upgrade, and Downgrade Instructions on page 166](#)
- [Product Compatibility on page 174](#)

## New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F7 for the MX Series.

- [Hardware on page 5](#)
- [Class of Service \(CoS\) on page 10](#)
- [Forwarding and Sampling on page 11](#)
- [General Routing on page 11](#)
- [High Availability and Resiliency on page 13](#)
- [Interfaces and Chassis on page 13](#)
- [IPv6 on page 17](#)
- [Management on page 17](#)
- [MPLS on page 17](#)
- [Multicast on page 18](#)
- [Network Management and Monitoring on page 18](#)
- [Routing Protocols on page 19](#)
- [Services Applications on page 22](#)
- [Software Installation and Upgrade on page 24](#)
- [Software-Defined Networking on page 24](#)
- [Subscriber Management and Services on page 27](#)
- [System Logging on page 34](#)
- [VPNs on page 35](#)

---

### Hardware

- **New Routing Engine RE-S-X6-64G (MX240, MX480, and MX960)**—Starting in Junos OS Release 15.1F4, the Routing Engine RE-S-X6-64G is supported on MX240, MX480, and MX960 routers. This Routing Engine has an increased computing capability and scalability to support the rapid rise in the data plane capacity. The Routing Engine is based on a modular virtualized architecture and leverages the hardware-assisted virtualization capabilities.



**NOTE:** Subscriber services and virtual-chassis support is not available in Junos OS 15.1Fx releases.

The Routing Engine has a 64-bit CPU and supports a 64-bit kernel and 64-bit applications. With its multicore capabilities, the Routing Engine supports symmetric multiprocessing in the Junos OS kernel and hosted applications.



**NOTE:** The Routing Engine RE-S-X6-64G is supported only on SCBE2, and it is not compatible with the SCB or the SCBE.

- **New rate-selectable MPC MPC7E-MRATE (MX2020, MX2010, MX960, MX480, and MX240)**—Starting in Junos OS Release 15.1F4, the rate-selectable MPC MPC7E (Multi-Rate) (model number: MPC7E-MRATE) is supported on MX2020, MX2010, MX960, MX480, and MX240 routers.

The main features of the MPC7E-MRATE MPC are the following:

- Line-rate throughput of up to 480 Gbps on MX240, MX480, and MX960 routers.
- Line-rate throughput of up to 400 Gbps on the MX2000 line of routers.
- Twelve ports that can each be configured as a 40-Gigabit Ethernet port or as four 10-Gigabit Ethernet ports by using a breakout cable. The ports support quad small-form factor pluggable plus (QSFP+) transceivers.
- Four ports—0/2, 0/5, 1/2, and 1/5—out of the twelve ports can be configured as 100-Gigabit Ethernet ports.
- You can configure different combinations of port speeds as long as the aggregate capacity per group of six ports labeled 0/0 through 0/5 does not exceed 240 Gbps. Similarly, aggregate capacity per group of the other six ports labeled 1/0 through 1/5 must not exceed 240 Gbps.



**NOTE:** To use the MPC7E-MRATE MPC on Junos OS Release 15.1F4, you must download and install the Junos Continuity software package for Junos OS Release 15.1F4.

Junos Continuity software package is not required for Junos OS Releases 15.1F6 and later.

- **Support for MPC8E on MX2010 and MX2020 routers**—Starting in Release 15.1F5, Junos OS supports MPC8E, a new Modular Port Concentrator (MPC) with two Modular Interface Card (MIC) slots, that provides a maximum bandwidth of 960 Gbps. MPC8E has four Packet Forwarding Engines, each providing a maximum bandwidth of 240 Gbps.



**NOTE:** To use the MPC8E MPC on Junos OS Release 15.1F5, you must download and install the Junos Continuity software package for Junos OS Release 15.1F5.

Junos Continuity software package is not required for Junos OS Releases 15.1F6 and later.

MPC8E supports:

- Line-rate throughput of up to 960 Gbps on the MX2000 line of routers.
- Two 12-port MIC-MRATE MICs with QSFP+ transceivers that support rate-selectability at the port level.
- Configuration of four ports out of 12 MIC-MRATE ports as 100-Gigabit Ethernet ports.
- Configuration of PIC-based tunnel interfaces from the Junos CLI, which allows you to configure 4,000 tunnel interfaces per PIC and 16,000 tunnel interfaces per line card.
- Maximum Transmission Unit (MTU) size of 16,000 bytes for transit traffic.
- [Dynamic power management](#) for effective utilization of available power.
- [Inline flow monitoring](#) for higher scalability and performance.
- [Flexible queuing](#) using an add-on license to support 32,000 queues per line card, including queues on both ingress and egress interfaces. You can use an additional license to support up to 512,000 queues per slot or 1,000,000 queues per slot.
- [Hyper mode](#) to speed up packet processing.
- **1-port 100-Gigabit DWDM OTN MIC with CFP2 (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Release 15.1F5, Junos OS supports the 1-port 100-Gigabit dense wavelength division multiplexing (DWDM) optical transport network (OTN) MIC (MIC3-100G-DWDM) with CFP2 analog coherent optical (CFP2-ACO) pluggable optics on MPC3E (MX-MPC3E-3D) and MPC3E NG (MPC3E-3D-NG). The 100-Gigabit Ethernet DWDM OTN MIC supports the following features:
  - Transparent transport of 100-Gigabit Ethernet signals with optical channel transport unit, OTU4 (V) framing.
  - Dual-polarization quadrature phase shift keying (DP-QPSK) modulation with coherent receiver and soft-decision forward error correction (SD-FEC) for long-haul and metro applications.
  - International Telecommunication Union (ITU)-standard OTN performance monitoring and alarm management
  - Extensive optical, digital signal processing (DSP), and bit error ratio (BER) performance monitoring statistics for the optical link.
- **New Routing Engine REMX2K-X8-64G (MX2010, MX2020)**—Starting in Junos OS Release 15.1F5 the Routing Engine REMX2K-X8-64G is supported on MX2010, and MX2020 routers. This Routing Engine has an increased computing capability and scalability to support the rapid rise in the data plane capacity. The Routing Engine is based on a modular virtualized architecture and leverages the hardware-assisted virtualization capabilities.
- **Support for MPC9E (MX2010 and MX2020)**—Starting with Junos OS Release 15.1F5, MX2020 and MX2010 routers support the Modular Port Concentrator (MPC) MPC9E (MX2K-MPC9E) with two Modular Interface Card (MIC) slots. MPC9E supports only the new 12-port rate-selectable MIC MIC-MRATE. MPC9E has four Packet Forwarding Engines, each with forwarding capability of up to 400 Gbps.



**NOTE:** To use MPC9E on Junos OS Release 15.1F5, you must download and install the Junos Continuity software package for Junos OS Release 15.1F5.

Junos Continuity software package is not required for Junos OS Release 15.1F6 and later.

MPC9E supports:

- Line-rate throughput of up to 1.6 Tbps on the MX2000 line of routers with SFB2.
- Two 12-port MIC-MRATE MICs with QSFP+ transceivers that support rate-selectability at the port level.
- Configuration of 8 ports out of the 12 MIC-MRATE ports as 100-Gigabit Ethernet ports.
- Configuration of PIC-based tunnel interfaces from the Junos CLI, which allows you to configure 4,000 tunnel interfaces per PIC and 16,000 tunnel interfaces per line card.
- Maximum Transmission Unit (MTU) size of 16,000 bytes for transit traffic.
- [Dynamic power management](#) for effective utilization of available power.
- [Inline flow monitoring](#) for higher scalability and performance.
- [Flexible queuing](#) using an add-on license to support 32,000 queues per line card, including queues on both ingress and egress interfaces. You can use an additional license to support up to 512,000 queues or one 1,000,000 queues per slot.
- [Hyper mode](#) to speed up packet processing.
- **Support for enhanced Switch Fabric Board (SFB2) for increased fabric bandwidth per slot (MX2010 and MX2020)**—Starting with Release 15.1F5, Junos OS supports an enhanced Switch Fabric Board (model number: MX2000-SFB2-S) that provides increased fabric bandwidth per slot. The MX2000 line of routers support SFB and SFB2, but not both at the same time. However, during an upgrade from SFB to SFB2, the MX2000 line of routers support both SFB and SFB2 at the same time for the duration of the upgrade.



**NOTE:** To use SFB2 on Junos OS Release 15.1F5, you must download and install the Junos Continuity software package for Junos OS Release 15.1F5.

Junos Continuity software package is not required for Junos OS Releases 15.1F6 and later.

- **Support for MPC7E 10G (MX2020, MX2010, MX960, MX480, and MX240)**—Starting with Junos OS Release 15.1F5, MX2020, MX2010, MX960, MX480, and MX240 routers support the Modular Port Concentrator (MPC) MPC7E 10G (MPC7E-10G). This is a fixed-configuration MPC with 40 10-Gigabit Ethernet ports. To use the MPC7E 10G on



Junos OS Release 15.1F5, you must download and install the Junos Continuity software package for Junos OS Release 15.1F5.

Junos Continuity software package is not required for Junos OS Releases 15.1F6 and later.



**NOTE:**

- On the MX2000 line of routers, the MPC7E 10G is plugged into an adapter card. Therefore, to use the MPC7E 10G MPC on MX2000 line of routers, the adapter card must be installed on the routers.
- To operate the MPC7E 10G on MX240, MX480, and MX960 routers, the routers must be equipped with high-capacity power supply, high-capacity fan tray, and Enhanced Switch Control Board SCBE2.

The main features of the MPC7E-10G MPC are the following:

- Line-rate throughput of up to 400 Gbps on MX240, MX480, MX960, MX2010, and MX2020 routers.
- Forty 10-Gigabit Ethernet ports. The ports support small-form factor pluggable plus (SFP+) transceivers.
- Supports maximum transmission units (MTUs) from 256 bytes through 16,000 bytes.
- Supports **hyper mode** to speed up packet processing.
- Supports **flexible queuing** by using an add-on license to support 32,000 queues per line card, including queues on both ingress and egress interfaces. You can use an additional license to support up to 512,000 queues.



**NOTE:** On MX240, MX480, and MX960 routers, the MPC7E 10G powers on only if the **network-services mode** on the router is configured as either **enhanced-ip** or **enhanced-ethernet**. On MX2000 routers, no additional configuration is required because by default the router operates in **enhanced-ip mode**.

- **Junos OS support for MX2008 routers**—Starting with Junos OS Release 15.1F7, Junos OS supports the MX2008 3D Universal Edge Router (model number: CHAS-MX2008). The MX2008 router is a 10-slot half-rack chassis with increased port density, but uses less space and consumes less power. Additionally, with the MX2008, you can scale bandwidth up to 1.6 Tbps per slot by using a chassis that is approximately half a rack in size.

The MX2008 router is an Ethernet-optimized edge router that provides both switching and carrier-class Ethernet routing. The router enables a wide range of business and residential applications and services, including high-speed transport and VPN services,

next-generation broadband multiplay services, and high-volume Internet data center internetworking.

- **Support for 12-port rate-selectable MIC (MIC-MRATE) on MPC8E and MPC9E (MX2010 and MX2020)**—Starting with Junos OS Release 15.1F5, the MPCs MPC8E and MPC9E support the 12-port rate-selectable MIC (MIC-MRATE) on the MX2000 line of routers. MIC-MRATE uses the quad small form-factor pluggable plus (QSFP+) transceiver for connectivity and supports port speeds of 100 Gbps, 40 Gbps, and 10 Gbps. MIC-MRATE also supports breakout cables, which you can use to split a 40-Gigabit Ethernet port into four 10-Gigabit Ethernet ports.

MIC-MRATE supports a maximum of 48 10-Gigabit Ethernet interfaces. On MPC8E with MIC-MRATE, you can configure four ports as 100-Gigabit Ethernet interfaces. On MPC9E with MIC-MRATE, you can configure eight ports as 100-Gigabit Ethernet interfaces.

MIC-MRATE also supports remote port identification. The LEDs for individual ports on the MIC and the related CLI commands help identify ports and assist in guided cabling.



**NOTE:** To use MIC-MRATE on Junos OS Release 15.1F5, you must download and install the Junos Continuity software package for Junos OS Release 15.1F5.

Junos Continuity software package is not required for Junos OS Releases 15.1F6 and later.

---

### Class of Service (CoS)

- **Copy ToS bits from incoming IP header to outer GRE IP header (MX Series with MPCs)**—Starting in Junos OS Release 15.1F5, you can set GRE tunnel interfaces to copy the ToS bits (DSCP value) from the incoming IPv4 header to the outer GRE IP header for transit traffic. You can set this at the individual GRE interface level by including the **copy-tos-to-outer-ip-header-transit** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level, or globally by including the **copy-tos-to-outer-service-type** (**[gre]** | **[mt]**) statement at the **[edit chassis]** hierarchy level.

You can also now rewrite the DSCP/IP precedence value in both the inner and outer headers with the **rewrite rules** (**[dscp]** | **[inet-precedence]**) **default protocol** (**[inet-both]** | **[inet-outer]**) statement at the **[edit class-of-service interfaces interface-name]** hierarchy level.

- **Support for packet marking schemes on a per-customer basis (MX Series only)**—Traditionally, packet marking in the Junos OS uses the forwarding class and loss priority determined from a BA classifier or multifield classifier. This approach does not allow for direct assignment of rewrite rules on a per-customer basis due to the limited number of combinations of forwarding class and loss priority.

Beginning with Junos OS release 15.1F6, there is a packet marking scheme, called policy map, that allows the definition of rewrite rules on a per-customer basis. Policy maps are defined at the **[edit class-of-service policy-map]** hierarchy level and can be assigned to a customer through a firewall action, an ingress interface, or a routing policy.

## Forwarding and Sampling

- **Enhancements to inline flow monitoring (MX Series)**—You can now configure Junos OS to set the flow record value to 0 (zero) for rejected, discarded, and policed flows. Starting with Junos OS release 15.1F7, you can configure the **report-zero-oif-gw-on-discard** statement at the **[edit chassis fpc slot-number inline-services]** hierarchy level to enable the FPC to set the value to 0 (zero) for the elements `ipNextHopIPv6Address` (Element ID 62), `egressInterface` (Element ID 14), and `ipNextHopIPv4Address` (Element ID 15) for the flow records for rejected, discarded, and policed flows.

If the **report-zero-oif-gw-on-discard** statement is not configured, the flow records display the available information for these elements, which is the default behavior.



**NOTE:** The **set chassis fpc slot-number inline-services report-zero-oif-gw-on-discard** command is not applicable for flows that are policed and sampled by the Packet Forwarding Engine for the egress interfaces.

After you configure the **report-zero-oif-gw-on-discard** statement, each sampled packet updates the forwarding action of that flow in the flow record. That is, the last sampled packet of the flow just before export determines the forwarding action of that flow. For example, in the case of a rate-limiting policer, forwarding action taken on a flow is not deterministic. The flow can be treated either as forwarded or as policed based on the forwarding action of the last sampled packet of that flow.

## General Routing

- **Support for virtualization on RE-S-X6-64G (MX240, MX480, MX960, MX2010, and MX2020)**—The Routing Engine RE-S-X6-64G supports virtualization for the following platforms:
  - MX240, MX480, and MX960—Junos OS Release 15.1F3 and later
  - MX2010 and MX2020—Junos OS Release 15.1F5S1 and later

Virtualization enables the router to support multiple instances of Junos OS and other operating systems on the same Routing Engine. One instance of Junos OS, which runs as a guest operating system, is launched by default. The user needs to log in to this instance for operations and management.

With virtualization of the Routing Engine, Junos OS supports new **request** and **show** commands associated with host and hypervisor processes. The commands are related to:

- Reboot, halt, and power management for the host
- Software upgrade for the host

- Disk snapshot for the host
- **Support for the combined operation of Synchronous Ethernet and Precision Time Protocol or hybrid mode (MX104)**—A combined operation of Synchronous Ethernet and Precision Time Protocol (PTP), also known as hybrid mode, is supported on the MX104 routers. In hybrid mode, the Synchronous Ethernet equipment clock (EEC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP (also known as IEEE 1588v2) for time synchronization.

Synchronous Ethernet and PTP provide frequency and phase synchronization; however, the accuracy in the order of nanoseconds is difficult to achieve through either 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 if 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.



**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.
  - Hybrid mode is not supported on integrated routing and bridging (IRB) and aggregated Ethernet interfaces configured on MX104 routers.
  - Unified in-service software upgrade (unified ISSU) is not supported when clock synchronization is configured for hybrid mode on MX104 routers.
- 
- **Support for PTP over IRB interfaces (MX104)**—Starting in Junos OS Release 15.1F5, MX104 routers support Precision Time Protocol (PTP) over integrated routing and bridging (IRB) interfaces. In releases before Junos OS Release 15.1F5, MX104 routers support PTP over physical Ethernet interfaces only.
  - **Enhancement to memory utilization (MX Series)**—Junos OS Release 15.1R5 supports an enhanced method for calculating the memory utilization by a Routing Engine. The inactive memory is now considered free and is no longer included in the calculation of memory utilization. That is, the value for used memory shown in the output of the **show**

**chassis routing-engine** command decreases and results in more memory to be available for other processes.

### High Availability and Resiliency

- **Support for unified ISSU on MX Series routers with MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Release 15.1F6, Junos OS supports unified in-service software upgrade (ISSU) on MX Series routers with MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q.

ISSU enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic.



**NOTE:** Unified ISSU is not supported on MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q with the following MICs:

- MS-MIC-16G
- MIC-3D-8DS3-E3
- MIC-3D-10C192-XFP

### Interfaces and Chassis

- **Support for targeted aggregated Ethernet distribution (MX Series routers with MPCs/MICs)**—In Junos OS Release 15.1F2 and later, you can direct traffic through specified links of a logical interface of an aggregate Ethernet bundle that is configured without link protection. This feature is supported on interfaces configured on MX Series MPCs and MICs.

By default, aggregated Ethernet bundles use a hash-based algorithm to distribute traffic over multiple links. Traffic destined through a logical interface of a bundle can exit through any of the member links based on the hashing algorithm. Therefore, egress policy enforcement might not always be accurate.

By configuring targeted aggregated Ethernet distribution, you can create distribution lists consisting of specific child member links. You can, therefore, enforce egress transit traffic to traverse through the specified links of the distribution lists. This configuration helps you enforce egress policies correctly. That is, you can implement policers on specific links that carry the desired traffic.



**NOTE:** Targeted aggregated Ethernet distribution can be applied to egress transit traffic only, excluding host outbound traffic.

- **Support for dynamic power management on MPC6E**—Starting in Junos OS Release 15.1F4, dynamic power management is supported on MPC6E on MX2010 and MX2020 routers. In earlier Junos OS releases, this feature is supported only on MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q on MX240, MX480, MX960, MX2010, and MX2020 3D Universal Edge Routers.

- **Support for flexible queuing on MPC5E**—Starting in Junos OS Release 15.1F4, flexible queuing is supported on MPC5E on MX240, MX480, MX960, MX2010, and MX2020 3D Universal Edge Routers.
- **Dynamic power management enabled by default**—Starting in Junos OS Release 15.1F4, dynamic power management is enabled by default. The **mic-aware-power-management** statement, which was used to enable dynamic power management in earlier releases, is deprecated.
- **Enhancement to ambient-temperature statement (MX Series)**—In Junos OS Release 15.1F4 and later, the default ambient temperature is set at 40° C on MX480, MX960, MX2010, and MX2020 3D Universal Edge Routers. You can override ambient temperature by setting the temperature at 55° C or 25° C.

```
[edit]
user@router# set chassis ambient-temperature ?
Possible completions:
25C                25 degree celsius
40C                40 degree celsius
55C                55 degree celsius
[edit]
```

When a router restarts, the system adjusts the power allocation or the provisioned power for the line cards on the basis of the configured ambient temperature. If enough power is not available, a minor chassis alarm is raised. However, the chassis continues to run with the configured ambient temperature. You can configure a new higher ambient temperature only after you make more power available by adding new power supply modules or by taking a few line cards offline. By using the provisioned power that is saved by configuring a lower ambient temperature, you can bring more hardware components online.

- **Support for centralized clocking on MX2008 routers**—Starting with Junos OS Release 15.1F7, the MX2008 router (model number: CHAS-MX2008) uses the centralized Stratum 3 clock module on the Routing and Control Board (RCB) to lock onto Synchronous Ethernet and distribute the frequency to the entire chassis. Supported features include:
  - Clock monitoring, filtering, and holdover
  - Hitless transition from a distributed to a centralized clocking mode
  - Distribution of the selected chassis clock source to downstream network elements by using supported line interfaces

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



**NOTE:** The MX2008 supports Precision Time Protocol (PTP) in distributed mode.

---

- **MPCs and MICs supported on MX2008 routers**—The MX2008 router (model number: CHAS-MX2008) supports all the MPCs (excluding AS-MCC) and MICs that are supported by the MX2000 line of routers.

MPCs native to the MX2000 line of routers (MPC6E, MPC8E, and MPC9E) are supported without an adapter card, but other MPCs (MS-MPC, MPC1, MPC2, MPC3, MPC4, MPC5, MPC7, MPC2E-NG, MPC3E-NG, and all variants) are supported with an adapter card.



**NOTE:** MX2008 routers do not support the Application Services Modular Carrier Card (AS-MCC).

[See [MPCs Supported by MX240, MX480, MX960, MX2010, and MX2020 Routers](#)]

- **Junos OS support for FRU management of MX2008 routers**—Starting with Junos OS Release 15.1F7, Junos OS supports the MX2008 router (model number: CHAS-MX2008). The Junos OS chassis management software for the MX2008 routers provides enhanced environmental monitoring and field-replaceable unit (FRU) control.

The MX2008 host subsystem consists of two Routing and Control Boards, or RCBs (model number REMX2008-X8-64G). The RCB is an integrated board and a single FRU that provides Routing Engine and Control Board functionality and supports virtualization. The router contains 8 SFBs (fabric cards, model number: MX2008-SFB2) that provides 7+1 redundancy. The router supports a maximum of 10 MPCs including adapter cards, and up to 20 MICS—a maximum of two MICs can be installed in each MPC.

The chassis contains nine power supply modules (PSMs) and two power distribution modules (PDMs) for the power feeds. Each PSM delivers 2500 W of power, and provides 8+1 redundancy. The two PDMs provide feed redundancy, with each PDM connected to primary and backup feeds separately.

The MX2008 cooling system contains two fan trays, with six fans in each. The fan trays can be installed at or removed from the back of the chassis, which allows the space in the front to be used for cable management. The MX2008 supports temperature thresholds for each temperature sensor, which enables the router to precisely control the cooling, raise alarms, and shut down a FRU.

- **Enhancement to policer configuration**—Starting in Junos OS Release 15.1F7, you can configure the MPC to take a value in the range 0 through 5 for the policer tick byte by using the **policer-limit** statement at the **[edit chassis]** hierarchy level. If this statement is not configured, the policer tick byte can take values till 7, which is the default behavior. You can use the **set chassis policer-limit** command to enable this feature.

You must restart the MPC or the router for the changes to take effect.

- **Software feature support on the MX2008**—Starting with Junos OS Release 15.1F7, the MX2008 router supports all software features that are supported by other MX Series routers in Junos OS Release 15.1F6.

The following key Junos OS features are supported:

- Basic Layer 2 features including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
- Class of service (CoS)
- Firewall filters and policers

- Integrated routing and bridging (IRB)
- Interoperability with existing MPCs (excluding the Application Services Modular Carrier Card, or AS-MCC)
- Layer 2 protocols
- Layer 2 VPNs, Layer 2 circuits, and Layer 3 VPNs
- Layer 3 routing protocols and MPLS
- Layer 3 services supported on MS-MIC and MS-MPC (for example, CGNAT, IP Security, inline active flow monitoring) and inline services
- Multicast forwarding
- Port mirroring
- Spanning-tree protocols, such as STP, MSTP, RSTP, and VSTP
- Synchronous Ethernet and Precision Time Protocol (IEEE 1588)
- Tunneling
- Graceful Routing Engine Switchover (GRES) and Non Stop Routing (NSR)



## IPv6

- **Forced IPv6 DNS server address insertion (MX Series)**—Starting in Junos OS Release 15.1F5, MX Series devices can dynamically provision DHCPv6 lease times and DNSv6 server IP addresses for DHCPv6 clients. The IP addresses and lease times are provided to DHCPv6 clients in DHCPv6 Advertisement and Reply messages without requiring a Solicit or Request message from a CPE device.

## Management

- **Junos Telemetry Interface enhancements (MX Series)** —Junos Telemetry Interface enables you to export telemetry data from supported interface hardware. Line card sensor data, such as physical interface events, are sent directly to configured collection points without involving polling. Starting with Junos OS Release 15.1F6, you can export LSP statistics and firewall filter statistics.

To enable the exporting of LSP statistics, include the **resource /junos/services/label-switched-path/usage/** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level. You must also configure the **sensor-based-stats** statement at the **[edit protocols mpls]** hierarchy level. Additionally, you must configure the **enhanced-ip** statement at the **[edit chassis network-services]** hierarchy level. Only dynamically configured LSPs and RSVP LSPs are supported. Statistics are not collected for P2MP LSPs, LDP LSPs, or static LSPs. To enable the exporting of data for firewall filters, include the **resource /junos/system/linecard/firewall/** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level.

Also starting with Junos OS Release 15.1F6, Junos Telemetry Interface is also supported on interfaces configured on the MPC7E, the MPC8E, and the MPC9E. Previously only MPC1 through MPC6E were supported.

## MPLS

- **Support for IS-IS segment routing (MX Series)**—Starting with Junos OS Release 15.1F5, IS-IS segment routing support is enabled through MPLS. Currently, label advertisements are supported for IS-IS only. IS-IS creates an adjacency segment per adjacency, per level, and per address family (one each for IPv4 and IPv6). Junos OS IS-IS implementation allocates node segment label blocks in accordance with the IS-IS protocol extensions for supporting segment routing node segments and provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:
  - **source-packet-routing**—Enable the source packet routing feature.
  - **node-segment**—Enable source packet routing at all levels.
  - **use-source-packet-routing**—Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and primary IS-IS source packet routing node segments.
  - **no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.

- **Subnet-match authentication for LDP sessions (MX Series)**—Starting with Junos OS Release 15.1F5, support for Hashed Message Authentication Code (HMAC) and MD5 authentication for LDP sessions is extended from a per-session configuration to a subnet-match (that is, longest-prefix-match) configuration.

This feature provides flexibility in configuring authentication for automatically targeted LDP (TLDP) sessions, making the deployment of remote loop-free alternate (LFA) and FEC 129 pseudowires easy.

To enable this feature, configure the **session-group** option at the **[edit protocols ldp session]** hierarchy level, and then enable the required authentication for the configured session group.

## Multicast

---

- **Protection against label spoofing or errant label injection across ASBRs (MX Series)**—Starting with Junos OS Release 15.1F2, you can use regular BGP implicit and explicit export policies to restrict VPN ASBR peer route advertisement to a given routing instance.

This is especially useful in the context of Inter-AS VPN Option-B ASBRs because it prevents a peer ASBR in a neighboring AS from spoofing or unintentionally injecting a VPN label intended for a different peer AS or intra-AS into the protected AS. In other words, service providers can configure a common ASBR so it does not accept MPLS packets from a peer ASBR unless the label has been explicitly advertised to the common ASBR.

Two new commands are introduced to provide this protection: **mpls-forwarding** at the **[edit routing-instances name instance-type mpls-forwarding]** hierarchy level and **forwarding-context** at the **[edit protocols bgp group group-name neighbor address]**, hierarchy level.

- **SAFI 129 NLRI compliance with RFC 6514 (MX Series)**—Starting with Junos OS Release 15.1F2, the NLRI format available for BGP VPN multicast is changing from the de facto format of SAFI 128 to SAFI 129 as defined in RFC 6514. SAFI 128 uses *length, label, prefix*. SAFI 129 uses *length, prefix*.

To use SAFI 129, enable the **rfc6514-compliant-safi129** statement at any of the following hierarchy levels: **[edit protocols bgp]**, **[edit protocols bgp group group-name]**, or **[edit protocols bgp group group-name neighbor address]**.

- **Latency fairness optimized multicast (MX Series)**—Starting with Junos OS Release 15.1F3, you can reduce latency in the multicast packet delivery by optimizing multicast packets sent to the Packet Forwarding Engines. You can achieve this by enabling the **ingress** or **local-latency-fairness** option in the **multicast-replication** configuration statement at the **[edit forwarding-options]** hierarchy level. The **multicast-replication** statement is supported only on platforms with the **enhanced-ip** mode enabled. This feature is not supported in VPLS networks and Layer 2 bridging.

## Network Management and Monitoring

---

- **SNMP support for fabric queue depth, WAN queue depth, and fabric counter (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Release 15.1F3, Junos OS

provides SNMP support for WAN queue depth, fabric queue depth, and fabric counter. The following SNMP MIB tables include the associated objects:

- **jnxCosQstatTable** table
- **jnxCosIngressQstatTable** table
- **jnxFabricMib** table

In addition, this feature supports the following traps for the Packet Forwarding Engine resource monitoring MIBs:

- **jnxPfeMemoryTrapVars**
- **jnxPfeMemoryNotifications**
- **Support for Timing MIB on MX104 router**—Starting in Junos OS Release 15.1F5, MX104 3D Universal Edge Router supports the timing feature. A new enterprise-specific MIB, Timing Feature Defect/Event Notification MIB, has been added to support this feature. The trap notifications are disabled by default. To enable SNMP trap notifications for timing events and defects, include the **timing-events** statement at the **[edit snmp trap-group trap-group object categories]** hierarchy level.

## Routing Protocols

- **Weighted ECMP support for one-hop IS-IS neighbors (MX Series)**—Beginning with Junos OS Release 15.1F4, you can configure the IS-IS protocol to get the logical interface bandwidth information associated with the gateways of equal-cost multipath (ECMP) next hop. During per-packet load balancing, traffic distribution is based on the available bandwidth to facilitate optimal bandwidth usage for incoming traffic on an ECMP path of one hop distance. The Packet Forwarding Engine does not distribute the traffic equally, but considers the balance values and distributes the traffic according to the bandwidth availability. However, this feature is not available for ECMP paths that are more than one hop away.
- **Support for BGP Optimal Route Reflection (BGP-ORR) (MX Series)**—Starting with Junos OS Release 15.1F4, you can configure BGP-ORR with IS-IS as the interior gateway protocol (IGP) on a route reflector to advertise the best path to the BGP-ORR client groups by using the shortest IGP metric from a client's perspective, instead of the route reflector's view.

To enable BGP-ORR, include the **optimal-route-reflection** statement at the **[edit protocols bgp group group-name]** hierarchy level.

Client groups sharing the same or similar IGP topology can be grouped as one BGP peer group. You can configure **optimal-route-reflection** to enable BGP-ORR in that BGP peer group. You can also configure one of the client nodes as the primary node (**igp-primary**) in a BGP peer group so that the IGP metric from that primary node is used to select the best path and advertise it to the clients in the same BGP peer group. Optionally, you can also select another client node as the backup node (**igp-backup**), which is used when the primary node (**igp-primary**) goes down or is unreachable.

Use the following CLI hierarchy to configure BGP-ORR:

```
[edit protocols bgp]
```

```
group group-name{
  optimal-route-reflection {
    igp-primary ipv4-address;
    igp-backup ipv4-address;
  }
}
```

Use the following CLI commands to monitor and troubleshoot the configuration for BGP-ORR:

- **show bgp group**—View the primary and backup configurations of BGP-ORR.
- **show isis bgp-orr**—View the IS-IS BGP-ORR metric (RIB).
- **show route advertising protocol bgp peer**—Verify whether the routes are being advertised according to the BGP-ORR rules.
- **IS-IS purge originator identification TLV (MX Series)**—Beginning with Release 15.1F4, Junos OS supports RFC 6232, *Purge Originator Identification TLV for IS-IS*, which defines a type, length and value (TLV) for identifying the origin of a purge initiated by the IS-IS protocol. You can configure this feature to add this TLV to a purge, along with the system ID of the Intermediate System (IS) that has initiated this purge. This makes it easier to locate the origin of the purge and its cause.
- **BGP flow specification for IPv6 (MX Series)**—Starting with Junos OS Release 15.1F5, this feature extends IPv6 support to the BGP flow specification which enables propagation of traffic flow specification rules for IPv6 and IPv6 VPN. The BGP flow specification automates coordination of traffic filtering rules in order to mitigate distributed denial-of-service attacks. In earlier Junos OS releases, flow-specific rules were propagated for IPv4 over BGP as network layer reachability information.  
  
To enable the BGP flow specification for IPv6, include the **flow** statement at the **[edit routing-options]** hierarchy level for global configuration or at the **[edit routing-instances routing-instance-name routing-options]** hierarchy level for instance-level configuration.
- **BGP labeled unicast supports stack of labels (MX Series)**—Beginning with Release 15.1F5, Junos OS supports RFC 3107, *Carrying Label Information in BGP-4*, that allows stacking of multiple labels in the BGP labeled unicast. In earlier Junos OS Releases, only one label per prefix was supported in the BGP unicast label. Junos OS now supports a label stack of up to five labels per prefix in the BGP labeled unicast updates. BGP labeled unicast updates with more than five labels are not supported, and Junos OS sets their state to **hidden**. This feature allows the use of BGP unicast label stack to control packet forwarding in the network and to reflect the BGP unicast label stack routes to its clients without changing the next hop.
- **Restricting LSP flooding over IS-IS interfaces (MX Series)**—Beginning with Junos OS Release 15.1F5, the IS-IS protocol can restrict flooding of LSAs to control sharing of routes between multiple level 2 metro ring networks. You can segregate both level 1 and level 2 networks into flood groups by using area IDs as tags to identify a flood group. Configure interfaces with specific area IDs to modify the flooding behavior as per your requirements. For example, when a router is connected to five level 2 metro ring networks, by default all the routers in the five rings are flooded with all LSP routes. You can configure five distinct flood groups on the ring-facing interfaces on the pre-aggregation device to restrict LSP flooding to a specific area. Configure area IDs

on interfaces to segregate them into flood groups. LSPs that belong to the specified area only are flooded through these interfaces. However, self-originated LSPs are not affected by this configuration.

- **Micro loop avoidance when IS-IS link fails (MX Series)**—Beginning with Release 15.1F5, Junos OS enables a device to defer IS-IS route download when an IS-IS link fails in order to avoid micro loops. When local links go down, the IS-IS protocol floods an entire area with the database. If the node connected to the local interface that has failed converges faster than the neighboring node, then the connected node redirects traffic to the converged path. This redirection can result in micro looping of traffic until the neighboring node converges. When the primary path of a protected node fails, the connected node does not need to converge quickly if the configured back up path is not impacted. In this case, traffic flow towards a converged path is deferred until the configured delay time.
- **System performance enhancements for rpd, Packet Forwarding Engine, and kernel (MX Series)**—Beginning with Junos OS Release 15.1F6, performance of the routing protocol process (rpd), the Packet Forwarding Engine, and the kernel is enhanced to speed up the process with which the rpd learns the route states and changes, and reflects these changes in the ASIC-based Packet Forwarding Engine residing in the line cards. The key enhancements are faster route download rates when a router comes up after a reboot, or when you add a new line card, and faster update of the data plane in convergence scenarios. We recommend disabling daemons, such as Layer 2 address learning process (l2ald) and connectivity-fault management process (cfmd) —if they are not required— to improve system performance. Though these enhancements are mainly for the MX Series, other platforms might see some performance improvements as well.

To maximize route download performance, increase the priority of the route-install job in the krt module of rpd. To increase the route-install job priority, configure the **dynamic-route-install-job-priority** statement at the **[edit routing-options forwarding-table]** hierarchy level. The **dynamic-route-install-job-priority** option is disabled by default. You can also specify the **threshold-length** and the **recover-length**.

- **threshold-length**—The priority of a job in the krt-queue is increased when the number of entries in the krt-queue exceeds this value. By default, the **threshold-length** is 50000.
- **recover-length**—The priority of a job in the krt-queue is restored to the default priority when the number of entries in the krt-queue falls below this value. By default, the **recover-length** is 45000.

The **dynamic-route-install-job-priority** configuration option is available in Junos OS 15.1F6 and later 15.1F releases only. Configuring the **dynamic-route-install-job-priority** option might not be required in future software releases because of system changes. Therefore, this option might not be available in Junos OS Release 16.1 and later releases.

## Services Applications

- **Support for inline LSQ logical interface**—Starting in Junos OS Release 15.1F4, MPC2E-3D-NG and MPC3E-3D-NG support inline LSQ logical interface when flexible queuing is enabled. The inline LSQ logical interface (referred to as lsq-) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC.
- **Class-of-service (CoS) marking and reclassification for the MS-MICs and MS-MPCs**—Starting with Junos Release 15.1F5, the MS-MIC and MS-MPC support CoS configuration, which enables you to configure Differentiated Services code point (DSCP) marking and forwarding-class assignment for packets transiting the MS-MIC or MS-MPC. You can configure the CoS service alongside the stateful firewall and NAT services, using a similar rule structure.

[See [Configuring CoS Rules](#).]

- **Exclude interfaces support in flowspec (rpd-infra) (MX Series)**—Starting in Release 15.1, Junos OS excludes applying the **flowspec** filter to traffic received on specific interfaces. A new term is added at the beginning of the **flowspec** filter that accepts any packet received on these specific interfaces. The new term is a variable that creates an exclusion list of terms attached to the forwarding table filter as a part of the flow specification filter.

To exclude the **flowspec** filter from being applied to traffic received on specific interfaces, you must first configure a **group-id** on such interfaces by including the family **inet** filter group **group-id** statement at the **[edit interfaces]** hierarchy level, and then attach the **flowspec** filter with the interface group by including the **flow interface-group group-id exclude** statement at the **[edit routing-options]** hierarchy level. You can configure only one **group-id** per routing instance with the **set routing-options flow interface-group group-id** statement.

- **Support for IKE and IPsec on NAPT-44 and NAT64 (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1F5, you can enable the passing of IKE and IPsec packets through NAPT-44 and NAT64 filters between IPsec peers that are not NAT-T compliant by using the IKE-ESP-TUNNEL-MODE-NAT-ALG on MS-MPCs and MS-MICs.

Use the following hierarchy to enable the IKE-ESP-TUNNEL-MODE-NAT-ALG:

```
[edit applications]
application ike-esp-application-name {
  application-protocol ike-esp-nat;
  protocol udp;
  destination-port 500;
  inactivity-timeout 3600;
}
application-set ike-esp-application-set-name {
  application ike-esp-application-name;
}

[edit services nat]
pool ike-isp-nat-pool-name {
  address ip-prefix;
```

```

    port automatic;
  }
  rule rule-name {
    match-direction input;
    term 0 {
      from {
        source-address address;
        application-sets ike-esp-application-set-name;
      }
      then {
        translated {
          source-pool ike-isp-nat-pool-name;
          translation-type napt-44;
        }
      }
    }
  }
}

```

- **Support for IP reassembly on an L2TP connection**—You can now configure the service interfaces on MX Series routers with MPC2E-3D-NG, MPC2E-3D-NG-Q, MPC3E-3D-NG, MPC3E-3D-NG-Q, and MPC5E to support IP packet reassembly on a Layer 2 Tunneling Protocol (L2TP) connection. The IP packet is fragmented over an L2TP connection when the packet size exceeds the maximum transmission unit (MTU) defined for the connection. Depending on the direction of the traffic flow, the fragmentation can occur either at the L2TP access concentrator (LAC) or at the L2TP network server (LNS), and reassembly occurs at the peer interface. (In an L2TP connection, a LAC is a peer interface for the LNS and vice versa.)

You can configure the service interfaces on the LAC or on the LNS to reassemble the fragmented packets before they can be further processed on the network. On a router running Junos OS, a service set is used to define the reassembly rules on the service interface. The service set is then assigned to the L2TP service at the **[edit services l2tp]** hierarchy level to configure IP reassembly for L2TP fragments.

You can view the reassembly statistics by using the **show services inline ip-reassembly statistics <fpc fpc-slot | pfe pfe-slot>** command.

See [IP Packet Fragment Reassembly for L2TP Overview](#)

- **Support for inline flow monitoring on MPC7E-MRATE, MPC8E, and MPC9E (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1F5, MPC7E-MRATE, MPC8E, and MPC9E support inline flow monitoring. Inline active flow monitoring provides for higher scalability and performance, as the scaling and performance are not dependent on the capacity of the services interface.
- **Exclude interfaces support in flowspec (rpd-infra) (MX Series)**—Starting in Release 15.1, Junos OS excludes applying the **flowspec** filter to traffic received on specific interfaces. A new term is added at the beginning of the **flowspec** filter that accepts any packet received on these specific interfaces. The new term is a variable that creates an exclusion list of terms attached to the forwarding table filter as a part of the flow specification filter.

To exclude the **flowspec** filter from being applied to traffic received on specific interfaces, you must first configure a **group-id** on such interfaces by including the family

**inet** filter group **group-id** statement at the **[edit interfaces]** hierarchy level, and then attach the **flowspec** filter with the interface group by including the **flow interface-group group-id exclude** statement at the **[edit routing-options]** hierarchy level. You can configure only one **group-id** per routing instance with the **set routing-options flow interface-group group-id** statement.

---

## Software Installation and Upgrade

- **Limited encryption Junos OS image (“Junos Limited”)** created for customers in **Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia (MX240, MX480, MX960, MX2010, MX2020)**—Starting in Junos OS Release 15.1F4, customers in the Eurasian Customs Union (currently comprising of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) should use the “Junos Limited” image for MX240, MX480, MX960, MX2010, and MX2020 routers instead of the “Junos Worldwide” image. The “Junos Limited” image does not have data-plane encryption and is intended only for countries in the Eurasian Customs Union because these countries have import restrictions on software containing data plane encryption. Unlike the “Junos Worldwide” image, the “Junos Limited” image supports control plane encryption through Secure Shell (SSH) and Secure Sockets Layer (SSL), thus allowing secure management of the system.



**NOTE:** The limited encryption Junos OS image (“Junos Limited”) is to be used by customers in Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia. Customers in all other countries should use the “Junos” image that was introduced in 15.1R1 to replace “Junos Domestic” image.

---

---

## Software-Defined Networking

- **Dynamic acquisition of network topology (MX Series)**—Starting in Junos OS Release 15.1F4, the network topology abstraction daemon (ntad) provides the functionality to dynamically acquire the network topology. The NorthStar Controller runs Junos OS in a virtual machine (VM) that uses BGP-LS (the preferred protocol) or OSPF/IS-IS to learn the network topology. In Junos OS, BGP-LS or IGP publishes the acquired topology it learns into the traffic engineering database, which provides an in-memory representation of the network topology. The network topology abstraction daemon produces a copy of the traffic engineering database that the topology server uses.
- **Standby and secondary LSPs (MX Series)**—Starting in Junos OS Release 15.1F4, standby and secondary LSPs provide an alternate route in the event the primary route fails. The tunnel ID, from node, to node, and IP address of a secondary or standby LSP are identical to that of the primary LSP. However, secondary and standby LSPs have the following differences:
  - A secondary LSP is not signaled until the primary LSP fails.
  - A standby LSP is signaled regardless of the status of the primary LSP.
- **PCC multiple template support (MX Series)**—Starting in Junos OS Release 15.1F4, you can create LSP templates to define a set of LSP attributes to apply to all PCE-initiated LSPs that provide a name match with the regular expression (regex) name specified in the template. By associating LSPs (through regex name matching)



with an LSP template, you can automatically enable or disable LSP attributes across any LSPs that provide a name match with the regex name.

- **PCC delegation of auto-bandwidth and TE++ (MX Series)**—Starting in Junos OS Release 15.1F4, a TE++ LSP includes a set of paths that are configured as a specific container statement and individual LSP statements, called sub-LSPs, which all have equal bandwidth. For TE++ LSPs, a normalization process resizes the LSP when either of the following two triggers occurs:

- A periodic timer occurs.
- Bandwidth thresholds are met.

These triggers elicit one of the following responses:

- No change is required.
- LSP splitting—add another LSP and distribute bandwidth across all the LSPs.
- LSP merging—delete an LSP and distribute bandwidth across all the LSPs.

For a TE++ LSP, the NorthStar Controller displays a single LSP with a set of paths. The LSP name is based on the matching prefix name of all members. The correlation between TE LSPs is based on association, and the LSP is deleted when there are no remaining TE LSPs.

- **IGP-based topology discovery (MX Series)**—Starting in Junos OS Release 15.1F4, the NorthStar Controller supports dynamic topology acquisition by using routing protocols (IS-IS, OSPF, and BGP LS) to obtain real-time topology updates.

- **Support of Internet draft draft-crabbe-pce-pce-initiated-lsp-03 for the stateful PCE-initiated LSP implementation (MX Series )**—In the partial client-side implementation of the stateful Path Computation Element (PCE) architecture, the implementation of PCE-controlled LSPs that are dynamically initiated by a PCE is currently based on version 1 of Internet draft draft-crabbe-pce-pce-initiated-lsp. Starting with Junos OS Release 14.2R4 and 15.1F4, this implementation is upgraded to support version 3, as defined in Internet draft draft-crabbe-pce-pce-initiated-lsp-03.

Releases earlier than Junos OS Release 14.2R4 support the older version of the PCE draft, causing interoperability issues between a Path Computation Client (PCC) running a previous release and a stateful PCE server that adheres to Internet draft draft-crabbe-pce-pce-initiated-lsp-03.

- **Support of Internet draft draft-ietf-pce-stateful-pce-07 for the stateful PCC implementation (MX Series )**—The partial client-side implementation of the stateful Path Computation Element (PCE) architecture is currently based on version 2 of Internet draft draft-ietf-pce-stateful-pce. Starting with Junos OS Release 14.2R4 and 15.1F4, this implementation is upgraded to support version 7, as defined in Internet draft draft-ietf-pce-stateful-pce-07.

Releases prior to 14.2R4 support the older version of the PCE draft, causing interoperability issues between a Path Computation Client (PCC) running a previous release and a stateful PCE server that adheres to Internet draft draft-ietf-pce-stateful-pce-07.

- **OVSDB support (MX80, MX240, MX480, MX960, MX2010, MX2020 routers)**—Starting with Junos OS Release 15.1F4, the Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which VMware NSX controllers and MX Series routers that support OVSDB can communicate. In an NSX multi-hypervisor environment, NSX controllers and MX routers can exchange control and statistical information via the OVSDB schema, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to entities in a physical network and the reverse.
- **PCEP-based discovery for P2MP LSPs (MX Series)**—Starting with Junos OS Release 15.1F6, Junos OS can be configured to send P2MP LSP information to a controller. The capability is enabled in the `[set protocols pcep]` hierarchy for either an individual PCE or a PCE group:  

```
set protocols pcep pce pce-id p2mp-lsp-report-capability
```

```
set protocols pcep group pce-group p2mp-lsp-report-capability
```
- **Support for TCP-MD5 as a mechanism for securing PCEP sessions (MX Series)**—Starting with Junos OS Release 15.1F6, the `authentication-key` and `authentication-key-chain` commands are available in the `set protocols pcep` hierarchy to secure sessions from the router to a controller through PCEP.

Use the following CLI command to bind an MD5 key to PCEP sessions:

```
set protocols pcep pce pce-id authentication-key MD5-key
```

Use the following CLI commands to bind a key chain to PCEP sessions:

```
set protocols pcep pce pce-id authentication-key-chain key-chain
```

```
set protocols pcep pce pce-id authentication-algorithm md5
```

In support of this feature, the output for the following `show` commands includes a new field, `pcep-session-auth`:

- `show protocols pcep`
- `show path-computation-client status`
- **Destination MAC address rewrites for OpenFlow (MX80, MX240, MX480, and MX960)**—Some types of network equipment that function as routers accept and handle packets only if the destination MAC address in the packet is the same as the MAC address of the Layer 3 interface on which the packet is received. To interoperate with these routers, connected devices must also be able to rewrite the destination MAC address of an incoming packet. Starting with Junos OS Release 15.1F6, an OpenFlow controller can configure an MX Series router that supports OpenFlow to rewrite the destination MAC address of an incoming packet.

[See [Understanding How the OpenFlow Destination MAC Address Rewrite Action Works.](#)]

## Subscriber Management and Services

- **New support for Framed-IP-Netmask for access-internal routes (MX Series)**—Starting in Junos OS Release 15.1F3, the mask value returned by RADIUS in the Framed-IP-Netmask attribute during PPP negotiation is considered for application to the access-internal route for the subscriber session. In earlier releases, the attribute mask is ignored and a /32 netmask is always applied, with the consequence that the address is set to the value of the Framed-IP-Address attribute returned by RADIUS.

Now, when the SDB\_FRAMED\_PROTOCOL attribute is equal to AUTHD\_FRAMED\_PROTOCOL\_PPP, the value of SDB\_USER\_IP\_MASK is set to 255.255.255.255 by default. This value is overridden by the Framed-IP-Netmask value, if present.

The IP Netmask field in the output of the **show subscribers** command now displays the default value of 255.255.255.255 or the actual value of Framed-IP-Netmask only when the SDB\_FRAMED\_PROTOCOL attribute is equal to AUTHD\_FRAMED\_PROTOCOL\_PPP.

- **Support for a static unnumbered interface with \$junos-routing-instance (MX Series)**—Starting in Junos OS Release 15.1F5, you can configure a static logical interface as the unnumbered interface in a dynamic profile that includes dynamic routing instance assignment by means of the **\$junos-routing-instance** predefined variable.



**NOTE:** This configuration fails commit if you also configure a preferred source address, either statically with the **preferred-source-address** statement or dynamically with the **\$junos-preferred-source-address** predefined variable.



**NOTE:** The static interface must belong to the routing instance, otherwise the profile instantiation fails.

In earlier releases, when the dynamic profile includes the **\$junos-routing-instance** predefined variable, you must do both of the following, else the commit fails:

- Use the **\$junos-loopback-interface-address** predefined variable to dynamically assign an address to the unnumbered interface. You cannot configure a static interface address.
- Use the **\$junos-preferred-source-address** predefined variable to dynamically assign a secondary IP address to the unnumbered interface. You cannot configure a static preferred source address.
- **Static provisioning of unique subscriber ID including interface description**—Starting in Junos OS Release 15.1F5, you can configure DHCP server and DHCP relay to concatenate the interface description with the username during the subscriber authentication or client authentication process. The interface description is separated from the other username fields by the specified delimiter, or by the default delimiter “.” if no delimiter is specified. The interface description can include either the logical interface description or the device interface description.



**NOTE:** Ensure that the specified delimiter is not part of the interface description.

Use the new **interface-description (device-interface | logical-interface)** configuration statement at one of the following hierarchy levels to specify that either the device interface description or logical interface description be concatenated to the other username fields:

- [edit forwarding-options dhcp-relay authentication username-include]
- [edit forwarding-options dhcp-relay dhcpv6 authentication username-include]
- [edit forwarding-options dhcp-relay dhcpv6 group *group-name* authentication username-include]
- [edit forwarding-options dhcp-relay group *group-name* authentication username-include]
- [edit logical-systems *logical-system-name* forwarding-options dhcp-relay authentication username-include]
- [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* authentication username-include]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication username-include]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* authentication username-include]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include]
- [edit logical-systems *logical-system-name* system services dhcp-local-server authentication username-include]
- [edit logical-systems *logical-system-name* system services dhcp-local-server group *group-name* authentication username-include]
- [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay authentication username-include]
- [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* authentication username-include]
- [edit routing-instances *routing-instance-name* system services dhcp-local-server authentication username-include]
- [edit routing-instances *routing-instance-name* system services dhcp-local-server group *group-name* authentication username-include]
- [edit system services dhcp-local-server authentication username-include]

- **[edit system services dhcp-local-server dhcpv6 authentication username-include]**
- **[edit system services dhcp-local-server dhcpv6 group group-name authentication username-include]**
- **[edit system services dhcp-local-server group group-name authentication username-include]**
- **Flat file output for service filter-based accounting**—Starting in Junos OS Release 15.1F5, you can configure service-based accounting to output to a flat file, in either IPDR or CSV format, as defined by the **accounting-options** configuration statement. To configure flat file output for service-based accounting, use the new **local flat-file-profile flat-file-profile-name** configuration statement at the **[edit access profile profile-name]** hierarchy level. Next, add the new **service-accounting** configuration statement at the **[edit accounting-options flat-file-profile flat-file-profile-name fields]** hierarchy level. Then either add the new **local** configuration statement at the **[edit access profile profile-name service accounting-order]** hierarchy level, or use the existing **activation-protocol** configuration statement at the **[edit access profile profile-name service accounting-order]** hierarchy level and activate the service through a CLI configuration or command.
- **Support for maximum session limits on L2TP service interfaces (MX Series)**—Starting in Junos OS Release 15.1F5, you can include the **l2tp-maximum-session number** statement at the **[edit interfaces service-interface]** hierarchy level to specify the maximum number of sessions that are allowed on an individual service interface (si). New session requests on an interface are accepted only when the session count is less than the maximum session limit. If the limit has been reached, subsequent requests are dropped and the LNS responds with a CDN message (Result Code 2, Error Code 4). When a pool of interfaces is configured, interfaces at the maximum limit are ignored in favor of an interface in the pool that has a lower session count.
- **Enhanced load balancing on L2TP physical service interfaces (MX Series)**—Starting in Junos OS Release 15.1F5, when a service interface in a service device pool is rebooted, session reconnects and new session requests are distributed based on the number of sessions on the available interfaces in the pool. The sessions are assigned to the interface with the fewest sessions. If more than one interface has the minimum number of sessions, then a random selection determines which interface gets the session.  
  
In earlier releases, session load balancing is a simple round-robin distribution among the interfaces. Consequently, fewer sessions are assigned to a newly rebooted interface than to the other interfaces. For example, consider a pool with two si interfaces, si-0/0/0 and si-1/0/0. Each has 100 sessions. If si-1/0/0 reboots, it drops all 100 sessions. As the sessions reconnect, they alternate between the two interfaces so that when all sessions have reconnected, si-0/0/0 has 150 sessions and the reconnected si-1/0/0 interface has only 50 sessions.  
  
Consider the same pool with the new behavior. As sessions reconnect, si-1/0/0 has fewer sessions (0 to start) than si-0/0/0 (100). Because the interface with the fewest sessions is selected, all sessions are assigned to si-1/0/0 until it reaches the same count as si-0/0/0.
- **Support for username stripping per routing instance (MX Series)**—Starting in Junos OS Release 15.1F5, you can configure a subscriber access profile so that a portion of

each subscriber login string is discarded and the remaining characters are used as a modified username by an external AAA server for session authentication and accounting. The modified username appears in RADIUS Access-Request, Acct-Start, and Acct-Stop messages; RADIUS-initiated disconnect requests; and change of authorization (CoA) requests. This username stripping configuration replaces a domain map configuration, but can be overridden by a AAA server.

Use the following statements at the **[edit access profile *profile-name* session-options strip-user-name]** hierarchy level to configure username stripping:

- **delimiter *delimiter***—Specify up to eight characters that the router uses to determine the boundary between the new modified username and the part of the original username that is discarded. There is no default delimiter.
- **parse-direction (*left-to-right* | *right-to-left*)**—Specify the direction in which the login string is examined until one of the configured delimiters is identified; **left-to-right** is the default. The delimiter and all characters to the right of the delimiter are discarded.

For example, consider a login string of **drgt21@example.com\$84** with the delimiters configured to be **/@\$%#**. If the parse direction is **left-to-right**, the **@** delimiter is reached first and the modified username is **drgt21**. If the parse direction is **right-to-left**, then the **\$** delimiter is reached first and the modified username is **drgt21@example.com**.



**BEST PRACTICE:** We recommend that you do not configure username stripping either when multiple user authentications are needed or when a global domain map is configured for the same subscribers covered by the AAA options configuration.

The **show network-access aaa subscribers session-id *id-number* detail** command displays the modified username in the Session Authentication Username field. The **clear network-access aaa subscriber username *username*** command requires you to specify the original, unstripped username (login string). The output of the **show subscribers** command displays the unstripped username, and when you issue the **show subscribers user-name *username*** command, you must specify the unstripped username.

- **AAA option sets to authorize and configure subscribers per routing instance to support username stripping (MX Series)**—Starting in Junos OS Release 15.1F5, you can include one or more of the following statements at the new **[edit access aaa-options *aaa-options-name*]** hierarchy level to define a set of AAA options for a subscriber or set of subscribers that username stripping is applied to:
  - **access-profile *profile-name***—Specify the name of the access profile that includes the username stripping configuration.
  - **aaa-context *aaa-context-name***—Specify the logical-system:routing-instance that the subscriber session uses for AAA (RADIUS) interactions like authenticating and accounting.
  - **subscriber-context *subscriber-context-name***—Specify the logical-system:routing-instance in which the subscriber interface is placed.



**NOTE:** Only the default (master) logical system is supported.

Use the **aaa-options *aaa-options-name*** statement at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit \$junos-interface-unit ppp-options]** hierarchy level to apply the attributes to PPP subscribers tunneled from the LAC to the LNS inline service interface.

Alternatively, use the **aaa-options *aaa-options-name*** statement at the **[edit access group-profile *profile-name* ppp-options]** hierarchy level to apply the attributes to PPP subscribers tunneled from LACs that are members of the user group.

Usernames are examined and modified according to the subscriber and AAA contexts specified in the option set. In the event of a conflict between option sets configured in both a group profile and a dynamic profile, the dynamic profile takes precedence.

- **Shared memory log supports filter-based debugging (MX Series)**—Starting in Junos OS Release 15.1F5, Junos OS supports filter-based debugging using the shared memory log.

Junos OS uses a shared memory space to store log entries for subscriber service daemons, such as jpppd, jdncpd, jl2tpd, autoconfd, bbe-smgd, authd, cosd, and dfwd. You can display the shared memory log (shmlog) output using the **show shmlog entries logname (*logname* | all) <filter *filter*> <flag-name *flag*>** command.

By default, shared memory logging is enabled. To disable the shmlog, at the **[edit system services subscriber-management]** hierarchy level, use the **set overrides shmlog disable** command.

By default, shmlog filtering is disabled. To enable shmlog filtering, at the **[edit system services subscriber-management overrides]** hierarchy level, use the **set shmlog filtering enable** command.

To display shmlog output for all daemon logs, use the **logname all** option with the **show shmlog entries** command. To limit shmlog output to a specific daemon log, provide the daemon name after the **logname** option followed by an asterisk. For example, **logname jpppd\*** or **logname authd\***.

To filter shmlog output, use the **filter *filter*** option with the **show shmlog entries logname (*logname* | all)** command. To display a list of valid filters, use the **show shmlog entries logname all ?** command.

You can also limit output to shmlog entries with specific flags, such as transmit-packets, configuration, sessionDb, and so on, using the **flag-name *flag*** option with the **show shmlog entries logname (*logname* | all)** command. To display a list of valid flags, use the **show shmlog entries logname all flag-name ?** command.

To direct shmlog output to a file, at the **[edit system services subscriber-management overrides]** hierarchy level, use the **set shmlog file <filename>** command. To view shmlog output stored in a text file, use the **show shmlog entries filename *filename*** command.

- **Configurable session limits for L2TP (MX Series)**—Starting in Junos OS Release 15.1F5, you can configure a limit on the maximum number of L2TP sessions allowed for the chassis, for all tunnels, for a tunnel-group, for a client group, and for a client. When the

session limit is reached, no new sessions can be established until the number of current sessions drops below the configured limit. These configured session limits have no effect on the maximum supported chassis limits that are imposed through the Juniper Networks license.

When an L2TP session request is initiated, the LNS checks whether the number of current active sessions is less than the configured limit in the following order: chassis > tunnel > tunnel group > session-limit group > client.

At each level, when the count is less or when no limit is configured, the check passes and the LNS proceeds to check the next level. If all levels pass the check, the session can be established. If at any level the current session count is equal to the configured limit, then the LNS rejects the session request and does not check any other level. When the LNS rejects a session request for an existing tunnel, it returns a Call-Disconnect-Notify (CDN) message with a result code and error code both set to 4 in response to the ICRQ. When the rejected request is for a new tunnel, the tunnel is established but the session fails to come up, causing the tunnel to come down because it has no sessions.

The LAC performs the same session limit check, but only for the chassis and tunnel levels. The LAC rejects requests by returning a PPP terminate message to the client.

Use the **maximum-sessions** statement at any of the following hierarchy levels:

- **[edit access profile *profile-name* client *client-name* l2tp]**
- **[edit services l2tp]**
- **[edit services l2tp sessions-limit-group]**
- **[edit services l2tp tunnel]**
- **[edit services l2tp tunnel-group *group-name*]**

Use the following commands to monitor the number of active sessions compared to the configured maximums:

- **show services l2tp client**—Display about all L2TP clients or a specific L2TP client.
- **show services l2tp session-limit-group**—Display information about all session-limit groups or a specific session limit group.
- **show services l2tp summary**—Display L2TP summary information, including sessions at the chassis level.
- **show services l2tp tunnel**—Display information about all L2TP tunnels or a specific L2TP tunnel.
- **show services l2tp tunnel-group**—Display information about all L2TP tunnel groups or a specific L2TP tunnel group.
- **Ensuring IPCP negotiation for IPv4 DNS addresses (MX Series)**—Starting in Junos OS Release 15.1F5, the router can prompt customer premises equipment (CPE) to negotiate both primary and secondary IPv4 DNS addresses during IPCP negotiation. This feature is useful when the CPE fails to send DNS address options in the IPCP configure request message, or when the options are sent but rejected. In earlier releases,



either situation results in no DNS address negotiation even though IPv4 DNS addresses are available on the router. This DNS option enables the router to control IPv4 DNS address provisioning for dynamic and static terminated PPPoE and LNS subscribers.

Specify the DNS negotiation option with the **ipcp-suggest-dns-option** statement at one of the following hierarchy levels:

- **[edit dynamic-profiles *profile-name* interfaces pp0 unit "\$junos-interface-unit" ppp-options]**—Configure the router to prompt the CPE to negotiate the DNS addresses for dynamic PPPoE subscribers.
- **[edit interfaces *interface-name* ppp-options]**—Configure the router to prompt the CPE to negotiate the DNS addresses for static PPPoE subscribers.
- **[edit dynamic-profiles *profile-name* interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]**—Configure the router to prompt the CPE to negotiate the DNS addresses for dynamic LNS subscribers.
- **[edit interfaces *si-slot/pic/port* unit *logical-unit-number* ppp-options]**—Configure the router to prompt the CPE to negotiate the DNS addresses for static LNS subscribers.
- **[edit access group-profile *profile-name* ppp-options]**—Configure the router to prompt the CPE to negotiate the DNS addresses for tunneled PPP subscribers with an LNS user group profile.
- **Dynamic subscriber and service management on statically configured interfaces (MX Series)**—Starting in Junos OS Release 15.1F5, enhanced subscriber management supports dynamic service activation and deactivation for static subscribers. These static subscribers work with the native Juniper Networks Session and Resource Control (SRC), or you can configure RADIUS to activate and deactivate the services with change of authorization (CoA) messages. Note, however, that with RADIUS, authentication failure does not prevent the underlying interface from coming up and forwarding traffic. Instead, it prevents the subscriber from coming up, and thus service activation/deactivation. Authorization parameters such as IP addresses, net masks, policy lists, and QoS are also not imposed when using RADIUS.

Use the following commands to provide administrative control of static subscribers:

- **request services static-subscribers login interface *interface-name***
- **request services static-subscribers logout interface *interface-name***
- **request services static-subscribers login group *group-name***
- **request services static-subscribers logout group *group-name***

Use the following commands to monitor static subscribers:

- **show static-subscribers**
- **show static-subscribers interface *interface-name***
- **show static-subscribers group *group-name***
- **New predefined variables and Juniper Networks VSAs for family any interface filters (MX Series)**—Starting in Junos OS Release 15.1F6, you can use the

\$junos-input-interface-filter and \$junos-output-interface-filter predefined variables to attach a filter to a dynamic interface created for family any. The filter names are derived from the Juniper Networks VSAs, Input-Interface-Filter (26-191) and Output-Interface-filter (26-192). These VSAs are conveyed in the following RADIUS messages: Access-Request, Acct-Start, Acct-Stop, and Acct-Interim-Interval. You can specify the variables as the filter names with **input** and **output** statements at the **[edit dynamic-profiles *profile-name* interfaces *interface-name* unit *logical-interface-number* filter]** hierarchy level.

- **New predefined variable to group subscribers on a physical interface (MX Series)**—Starting in Junos OS Release 15.1F6, you can specify the new Juniper Networks predefined variable, \$junos-phy-ifd-interface-set-name, with the **interface-set** statement at the **[edit dynamic-profiles *profile-name* interfaces]** hierarchy level to configure an interface set associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

Another use case is optimizing CoS level 2 node resources by grouping residential subscribers into an interface set associated with the physical interface in a topology where residential and business subscribers share the interface, enabling the use of CoS level 2 nodes for the interface set rather than for each residential interface.

- **Configuring default values for routing instances (MX Series)**—Starting in Junos OS Release 15.1F6, you can define a default value for the Juniper Networks predefined variable, \$junos-routing-instance. This value is used in the event RADIUS does not supply a value for \$junos-routing-instance. To configure a default value, use the **predefined-variable-defaults** statement at the **[edit dynamic-profiles]** hierarchy level. For example, to set the default value to RI-default:

```
[edit dynamic-profiles profile-name]  
user@host# set predefined-variable-defaults routing-instance RI-default
```

- **DHCP rate adjustment (MX Series)**—Starting in Junos OS Release 15.1F7, you can use DHCP tags to modify the CLI-configured and RADIUS-configured shaping rate values after a subscriber is instantiated. The new values are conveyed in DHCP option 82, suboption 9 discovery packets. Suboption 9 contains the Internet Assigned Numbers Authority (IANA) DSL Forum VSA (vendor ID 3561).

Configure the shaping rate adjustment controls by including the **dhcp-tags** statement at the **[edit class-of-service adjustment-control-profiles *profile-name* application]** hierarchy level. Specify the desired rate-adjustment algorithm and set a priority for the DHCP Tags application in the adjustment control profile.

---

## System Logging

- **System log messages to indicate checksum errors on the DDR3 interface**—Starting in Junos OS Release 13.3 R9, two new system log messages, XMCHIP\_CMERROR\_DDRIF\_INT\_REG\_CHKSUM\_ERR\_MINOR and XMCHIP\_CMERROR\_DDRIF\_INT\_REG\_CHKSUM\_ERR\_MAJOR, are added to indicate memory-related problems on the interfaces to the double data rate type 3 (DDR3) memory. These error messages indicate that an FPC has detected a checksum error, which is causing packet drops.

The following error threshold values classify the error as a major error or a minor error:

- Minor error— 6 to 254 errors per second
- Major error—255 and more errors per second

## VPNs

- **VPLS dynamic profiles not supported with 64-bit rpd (MX Series)**— Starting with Junos OS Release 15.1F3, virtual private LAN service (VPLS) dynamic profiles are not supported with the 64-bit mode routing protocol process (rpd). A new system log error (**RPD\_DYN\_CFG\_64RPD\_UNSUPPORTED**) is displayed when this condition occurs indicating that rpd failed to notify the dynamic configuration clients about its availability to process the dynamic configuration requests. To enable the VPLS dynamic profiles configuration and use 32-bit mode, configure rpd by using the **set system process routing force-32-bit** command in the CLI.
- **Ethernet VPN multihoming—Ethernet segment identifier per interface (MX Series)**—Starting in Junos OS Release 15.1F6, Junos OS enables the Ethernet VPN (EVPN) multihoming feature to connect a customer site to two or more PE devices to provide redundant connectivity. A CE device can be multihomed to different PE devices or the same PE device. A redundant PE device can provide network service to the customer site as soon as a failure is detected. EVPN multihoming helps to maintain EVPN service and traffic forwarding to and from the multihomed site if one of the following types of network failure occurs:
  - PE device to CE device link failure
  - PE device failure
  - MPLS-reachability failure between the local PE device and a remote PE device

- See Also**
- *Changes in Behavior and Syntax*
  - *Known Behavior*
  - *Known Issues*
  - *Resolved Issues*
  - *Documentation Updates*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Product Compatibility*

## Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 15.1F7 for the MX Series..

- [Authentication, Authorization and Accounting \(AAA\) \(RADIUS\) on page 37](#)
- [Flow-based and Packet-based Processing on page 37](#)

- [General Routing on page 37](#)
- [High Availability \(HA\) and Resiliency on page 37](#)
- [Interfaces and Chassis on page 37](#)
- [IPv6 on page 38](#)
- [MPLS on page 38](#)
- [Network Management and Monitoring on page 38](#)
- [Routing Policy and Firewall Filters on page 39](#)
- [Routing Protocols on page 39](#)
- [Services Applications on page 40](#)
- [Subscriber Management and Services \(MX Series\) on page 41](#)
- [System Logging on page 43](#)
- [System Management on page 49](#)
- [Platform and Infrastructure on page 49](#)
- [Virtual Chassis on page 51](#)
- [VPNs on page 51](#)

### Authentication, Authorization and Accounting (AAA) (RADIUS)

---

- **Statement introduced to enforce strict authorization**—Starting in Junos OS Release 15.1F2, customers can use the **set system tacplus-options strict-authorization** statement to enforce strict authorization to the users. When a user is logging in, Junos OS issues two TACACS+ requests—first is the authentication request and then the authorization request. By default, when the authorization request is rejected by the TACACS+ server, Junos OS disregards this rejection and provides full access to the user. When the **set system tacplus-options strict-authorization** statement is set, Junos OS denies access to the user on failure of the authorization request.

### Flow-based and Packet-based Processing

---

- **Correction for Inline Jflow reporting (MX Series)**—Starting in Junos OS Release 15.1F2, when a destination is reachable through multiple paths, Inline Jflow reports OIF, GW, DST\_MASK, and DST\_AS data incorrectly in flow records. The new configuration statement **set services flow-monitoring <version-ipfix | version9> template <template\_name> nexthop-learning enable** corrects OIF, GW, DST\_MASK, and DST\_AS data reporting.

### General Routing

---

- **Modified output of the clear services sessions | display xml command (MX Series)**—In Junos OS Release 14.1X55-D30, the output of the **clear services sessions | display xml** command is modified to include the **<sess-marked-for-deletion>** tag instead of the **<sess-removed>** tag. In releases before Junos OS Release 14.1X55-D30, the output of this command includes the **<sess-removed>** tag. The replacement of the **<sess-removed>** tag with the **<sess-marked-for-deletion>** tag aims at establishing consistency with the output of the **clear services sessions** command that includes the field **Sessions marked for deletion**.

### High Availability (HA) and Resiliency

---

- **New ISSU warning message for VCCV-BFD NSR not being supported**—Starting in Junos OS Release 15.1R2, 15.1F2, and later releases, the Junos OS CLI displays a warning message (when you perform a unified in-service software upgrade (ISSU)) about NSR not being supported for Bidirectional Forwarding Detection (BFD) support for virtual circuit connectivity verification (VCCV). You must enter a **[yes]** or **[no]** input to confirm whether you want to proceed with the unified ISSU operation or not.

### Interfaces and Chassis

---

- **Support for automatic enabling of flow control for MACsec (MX Series)**—Starting in Junos OS Release 15.1F6, when Media Access Control Security (MACsec) is enabled on an interface, the interface flow control capability is enabled by default, regardless of the configuration that you set using the **(flow-control | no-flow-control)** statement at the **[edit interfaces interface-name gigether-options]** hierarchy level. When MACsec is disabled, interface flow control is restored to the configuration that you set using the **flow-control** statement at the **[edit interfaces]** hierarchy level. When MACsec is enabled, additional header bytes are added to the packet by the MACsec PHY. With line rate traffic, when MACsec is enabled and flow control is disabled, the pause frames sent by the MACsec PHY are terminated by the MIC's MAC (enhanced 20-port Gigabit Ethernet MICs on MX Series routers) and not transferred to the Packet Forwarding Engine, causing framing errors. Therefore, when MACsec is enabled on an interface, flow control is also automatically enabled on such an interface.

---

## IPv6

- **IPv6 addresses with padded zeros in MIC or MS-MPC system log messages (MX Series)**—Starting with Junos OS Release 15.1F5, all system log messages originating from MIC or MS-MPC line cards display padded zeros in IPv6 addresses to make them compatible with MS-DPC line cards. Earlier, the system log messages from MIC or MS-MPC line cards displayed IPv6 addresses with ':' instead of padded zeros.

---

## MPLS

- **Inline BFD support on IRB interfaces (MX Series routers with MPCs or MICs)**—Starting with Junos OS Release 15.1F4, the inline BFD sessions transmitted or received from FPC hardware are supported on integrated routing and bridging (IRB) interfaces. This enhancement is available only on MX Series routers with MPCs/MICs that have configured the **enhanced-ip** option.

---

## Network Management and Monitoring

- **New 64-bit counter of octets for interfaces (MX Series)**—Starting with Release 15.1F4, Junos OS supports two new Juniper Networks enterprise-specific Interface MIB Extension objects—**ifHCIn1SecOctets** and **ifHCOut1SecOctets**—that act as 64-bit counters of octets passing through an interface.
- **Enhancement for SONET interval counter (MX Series)**—Starting with Junos OS Release 15.1F5, only the **Current Day Interval Total** output field in the **show interfaces interval** command for SONET interfaces is reset after 24 hours. In addition, the **Previous Day Interval Total** output field displays the last updated time in **hh:mm**.  
[See [show interfaces interval](#).]
- **SNMP proxy feature (MX Series)**—Starting with Junos OS Release 15.1F2, you must configure **interface <interface-name>** statement at the **[edit snmp]** hierarchy level for the proxy SNMP agent. Earlier, configuring interface for the proxy SNMP agent was not mandatory.

## Routing Policy and Firewall Filters

---

- **Support for logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol (MX Series)**— Starting with Junos OS Release 15.1F6, you can configure logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol. The queue-depth indicates the number of IP options packets which can be enqueued in the Packet Forwarding Engine logical queue, beyond which it would start dropping the packets.

## Routing Protocols

---

- **Support for RFC 5492, *Capabilities Advertisement with BGP-4***—Beginning with Junos OS Release 15.1F5, BGP sessions can be established with legacy peers that do not support optional parameters, such as capabilities. In earlier Junos OS releases from 15.1R1 through 15.1R3 and 15.1F1 through 15.1F4, BGP sessions with legacy routers without BGP capabilities were not supported. Starting with Junos OS Release 15.1F5, support for BGP sessions with legacy routers without BGP capabilities is restored.
- **Timers of delay-route-advertisements are modified**—Beginning with Junos OS Release 15.1F7, the range of the timer values of **delay-route-advertisements** has been increased to 36000 from 3600. The default value of **route age**, that is the maximum delay after route aggregates have been created has also been modified to 0. In earlier Junos releases, the default **route age** was 1200. The timer values of **delay-route-advertisements** are configured to avoid premature route advertisements that might result in traffic loss in a BGP session.

[See [delay-route-advertisements](#).]

- **Change in default behavior of router capability (MX Series)**—In Junos OS Releases 15.1F7, 16.1R4, 16.1X65, and 17.1R1 and later releases, router capability TLV distribution flag (S-bit), that controls IS-IS advertisements, will be reset, so that the segment routing capable sub-TLV is propagated throughout the IS-IS level and not advertised across IS-IS level boundaries.
- **New option to delay BGP route advertisements (MX Series)**—Beginning with Junos OS Release 15.1F6, you can delay BGP route updates to its peers until the forwarding table is synchronized. This is to avoid premature route advertisements that might result in traffic loss. A new configuration statement **delay-route-advertisements** is available at the **[edit routing-instances routing-instance-name protocols bgp group group-name family inet unicast]** hierarchy level. You can configure both minimum and maximum delay periods to suit your network requirements.

[See [delay-route-advertisements](#).]

## Services Applications

---

- **Anycast address 0/0 must not be accepted in the from-clause of Detnat rule (MX Series)**—Starting with Junos OS release 15.1F5, for multiservices (ms-) interfaces, anycast configuration is not allowed as the source-address when translation type is deterministic NAT.
- **Change to show services nat pool command output**—Starting in Junos OS Release 15.1F5, the **show services nat pool** command output includes this new field: AP-P port limit allocation errors. When AP-P is configured, this field indicates the number of out-of-port errors that are due to a configured limit for the number of allocated ports in the **limit-ports-per-address** statement at the **[edit services nat pool nat-pool-name]** hierarchy level.
- **Disabling NAT-traversal for IPsec-protected packets (MX Series)**—Starting in Junos OS release 15.1F6, you can include the **disable-natt** statement at the **[edit services ipsec-vpn]** hierarchy level to disable NAT-traversal (NAT-T) on MX Series routers. When you disable NAT-T, the NAT-T functionality is globally switched off. Also, even when a NAT device is present between the two IPsec gateways, only Encapsulating Security Payload (ESP) is used when you disable NAT-T. When NAT-T is configured, IPsec traffic is encapsulated using the UDP header and port information provided for the NAT devices. By default, Junos OS detects whether either one of the IPsec tunnels is behind a NAT device and automatically switches to using NAT-T for the protected traffic. However, in certain cases, NAT-T support on MX Series routers might not work as desired. Also, you might require NAT-traversal to be disabled if you are aware that the network uses IPsec-aware NAT. In such cases, you can disable NAT-T.
- **Forwarding class and DSCP configuration for sampled packets (MX Series)**—Starting with Junos Release OS 15.1F6, you can configure the forwarding class and the Differentiated Services Code Point (DSCP) mapping that is applied to exported packets for inline active flow monitoring. Configure **forwarding-class class-name** and **dscp dscp-value** at the **[edit forwarding-options sampling instance instance-name family (inet | inet6) output flow-server hostname]** hierarchy level.

The *dscp-value* range is 0 through 63 (the default is 0). When the same **flow-server** is configured under both the **inet** and **inet6** families in a sampling instance, use the same **dscp** value for both **flow-server** appearances.

The *dscp-value* is overwritten by the CoS DSCP value if you configure **dscp** at the **[edit class-of-service]** hierarchy level.
- **Change in flow table size configuration**—Starting in Junos OS Release 15.1F2, the default value for the **ipv4-flow-table-size** is 1024, and changing the value for the **ipv4-flow-table-size** or **ipv6-flow-table-size** statement does *not* initiate an automatic reboot of the FPC.
- **Deprecated security idp statements (MX Series)**—The **[edit security idp]** configuration statements are deprecated for the MX Series for Junos OS Release 15.1F7.



## Subscriber Management and Services (MX Series)

- **Local DNS configurations available when authentication order is set to none (MX Series)**—Starting in Junos OS Release 15.1F2, subscribers get the DNS server addresses when both of the following are true:
  - The authentication order is set to **none** at the **[edit access profile profile-name authentication-order]** hierarchy level.
  - A DNS server address is configured locally in the access profile with the **domain-name-server**, **domain-name-server-inet**, or **domain-name-server-inet6** statement at the **[edit access profile profile-name]** hierarchy level.

In earlier releases, subscribers get an IP address in this situation, but not the DNS server addresses.

- **Support for longer CHAP challenge local names (MX Series)**—Starting in Junos OS Release 15.1F4, the supported length of the CHAP local name is increased to 32 characters. In earlier releases, only eight characters are supported even though the CLI allows you to enter a longer name. You can configure the name with the **local-name** statement at the **[edit dynamic-profiles profile-name interfaces pp0 unit "\$junos-interface-unit" ppp-options]** or **[edit dynamic-profiles profile-name interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]** hierarchy levels. The maximum length of the local name for PAP authentication remains unchanged at eight characters.
- **Increased maximum limits for accounting and authentication retries and timeouts (MX Series)**—Starting in Junos OS Release 15.1F5, you can configure a maximum of 100 retry attempts for RADIUS accounting (**accounting-retry** statement) or authentication (**retry** statement). In earlier releases, the maximum value is 30 retries. You can also configure a maximum timeout of 1000 seconds for RADIUS accounting (**accounting-timeout** statement) or authentication (**timeout** statement). In earlier releases, the maximum timeout is 90 seconds.



**NOTE:** The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

- **Extended range for RADIUS request rate (MX Series)**—Starting in Junos OS Release 15.1F6, the range for the **request-rate** statement at the **[edit access radius-options]** hierarchy level has been extended to 100 through 4000 requests per second. In earlier releases, the range is 500 through 4000 requests per second. The default value is unchanged at 500 requests per second.
- **Change in PPP keepalive interval for inline services subscribers (MX Series)**—Starting in Junos OS Release 15.1F7, you can configure the PPP keepalive interval for subscriber services in the range 1 second through 600 seconds. Subscriber PPP keepalives are handled by the Packet Forwarding Engine. If you configure a value greater than 600 seconds, the number is accepted by the CLI, but the Packet Forwarding Engine limits the interval to 600 seconds. The interval is configured in a PPP dynamic profile with

the **interval** statement at the **[edit dynamic-profiles *profile-name* interfaces pp0 unit \$junos-interface-unit keepalives]** hierarchy level.

In earlier Junos OS releases the range is from 1 second through 60 seconds. The Packet Forwarding Engine limits any higher configured value to an interval of 60 seconds.

PPP keepalives for non-subscriber services are handled by the Routing Engine with an interval range from 1 second through 32,767 seconds.

## System Logging

---

- **New JSERVICES system log messages (MX Series)**—In Junos OS Release 15.1F6, you can configure MX Series routers with MS-MPCs to log the following messages:

**Table 1: JSERVICES System Logs**

Name	System Log Message	Description	Severity
JSERVICES_ALG_FTP_ACTIVE_ACCEPT	software-string <i>src-ip:src-port</i> [ <i>xlated-src-ip:xlated-src-port</i> ]->[ <i>xlated-dst-ip:</i> <i>xlated-dst-port</i> ] <i>dst-ip:dst-port (protocol-name)</i>	An FTP data connection from client to server is established. The matching packet contains the indicated information about its protocol name, application, source (logical interface name, IP address, and port number), and destination (IP address and port number). If the flow requires NAT services, NAT information appears in the message.	LOG_NOTICE
JSERVICES_ALG_FTP_PASSIVE_ACCEPT	software-string <i>src-ip:src-port</i> [ <i>xlated-src-ip:xlated-src-port</i> ]->[ <i>xlated-dst-ip:</i> <i>xlated-dst-port</i> ] <i>dst-ip:dst-port (protocol-name)</i>	An FTP data connection from server to client is established. The matching packet contains the indicated information about its protocol name, application, source (logical interface name, IP address, and port number), and destination (IP address and port number). If the flow requires NAT services, NAT information appears in the message.	LOG_NOTICE
JSERVICES_DROP_FLOW_DELETE	software-string <i>src-ip:src-port</i> [ <i>xlated-src-ip:xlated-src-port</i> ]->[ <i>xlated-dst-ip:</i> <i>xlated-dst-port</i> ] <i>dst-ip:dst-port (protocol-name)</i>	The session with the indicated characteristics is removed and it had drop flow. The NAT data is available in the message if the session requires NAT.	LOG_NOTICE

Table 1: JSERVICES System Logs (continued)

JSERVICES_ICMP_ERROR_DROP	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The ICMP error packet was dropped because it did not belong to an existing flow.	LOG_NOTICE
JSERVICES_ICMP_HEADER_LEN_ERROR	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The ICMP packet was discarded because the length field in the packet header was shorter than the minimum 8 bytes required for an ICMP packet.	LOG_NOTICE
JSERVICES_ICMP_PACKET_ERROR_LENGTH	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The ICMP packet was discarded because the packet contained fewer than 48 bytes or more than 576 bytes of data.	LOG_NOTICE
JSERVICES_IP_FRAG_ASSEMBLY_TIMEOUT	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet and all related IP fragments previously received were discarded because all fragments did not arrive within the reassembly timeout period of 4 seconds.	LOG_NOTICE
JSERVICES_IP_FRAG_OVERLAP	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because the contents of two fragments overlapped.	LOG_NOTICE
JSERVICES_IP_PACKET_CHECKSUM_ERROR	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because checksum was incorrect.	LOG_NOTICE
JSERVICES_IP_PACKET_DST_BAD	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its destination address was either a multicast address or was in the range reserved for experimental use (248.0.0.0 through 255.255.255.254).	LOG_NOTICE

Table 1: JSERVICES System Logs (continued)

JSERVICES_IP_PACKET_FRAG_LEN_INV	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because the length of a fragment was invalid.	LOG_NOTICE
JSERVICES_IP_PACKET_INCORRECT_LEN	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The IP packet is discarded because packet length was invalid.	LOG_NOTICE
JSERVICES_IP_PACKET_LAND_ATTACK	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its source and destination address for the packet were the same (referred as land attack).	LOG_NOTICE
JSERVICES_IP_PACKET_LAND_PORT_ATTACK	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its source and destination address for the packet were the same and also its source and destination ports were same (referred as land port attack).	LOG_NOTICE
JSERVICES_IP_PACKET_NOT_VERSION_4	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet version was not IPv4.	LOG_NOTICE
JSERVICES_IP_PACKET_NOT_VERSION_6	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet version was not IPv6.	LOG_NOTICE
JSERVICES_IP_PACKET_PROTOCOL_ERROR	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because it used invalid IP protocol.	LOG_NOTICE
JSERVICES_IP_PACKET_SRC_BAD	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its source address was one of the following: (1) a multicast address (2) a broadcast address (3) in the range 248.0.0.0 through 255.255.255.254, which is reserved for experimental use.	LOG_NOTICE

Table 1: JSERVICES System Logs (continued)

JSERVICES_IP_PACKET_TTL_ERROR	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet with the indicated characteristics is discarded because the packet had a time-to-live (TTL) value of zero.	LOG_NOTICE
JSERVICES_IP_PACKET_TOO_LONG	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because the packet contained more than 64 kilobytes (KB) of data (referred to as a ping-of-death attack).	LOG_NOTICE
JSERVICES_IP_PACKET_TOO_SHORT	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet did not contain the minimum amount of data required.	LOG_NOTICE
JSERVICES_NO_IP_PACKET	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	Packet received was not an IPv4 or IPv6 packet.	LOG_NOTICE
JSERVICES_SYN_DEFENSE	proto <i>protocol-id</i> ( <i>protocol-name</i> ), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet with the indicated characteristics was discarded because the TCP handshake that is used to establish a session did not complete within the set time limit. The time limit is set by the 'open-timeout' statement at the [edit interfaces <services-interface> services-options] hierarchy level. If the time limit is not set, the session uses the default timeout value.	LOG_NOTICE
JSERVICES_SFW_NO_POLICY	source-ip:destination-ip No policy	The stateful firewall received packets with the indicated source and destination addresses. There was no matching policy for the traffic.	LOG_NOTICE

Table 1: JSERVICES System Logs (continued)

JSERVICES_SFW_NO_RULE_DROP	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The stateful firewall discarded the packet with the indicated characteristics, because the packet did not match any stateful firewall rules. In this case, the default action is to discard the packet. The discarded packet contained the indicated information about its protocol (numerical identifier and name), source (logical interface name, IP address, and port number), and destination (IP address and port number).	LOG_NOTICE
JSERVICES_TCP_FLAGS_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The packet was discarded because the flags in the packet were set in one of the following combinations: (1) FIN and RST (2) SYN and one or more of FIN, RST, and URG.	LOG_NOTICE
JSERVICES_TCP_HEADER_LEN_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The packet was discarded because the length field in the packet header was shorter than the minimum 20 bytes required for a TCP packet.	LOG_NOTICE
JSERVICES_TCP_NON_SYN_FIRST_PACKET	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The TCP packet was discarded because it was the first packet in the TCP session but the SYN flag was not set.	LOG_NOTICE
JSERVICES_TCP_PORT_ZERO	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The packet was discarded because the source or destination port specified in the packet was zero.	LOG_NOTICE



Table 1: JSERVICES System Logs (continued)

JSERVICES_TCP_SEQNUM_AND_FLAGS_ZERO	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The packet was discarded because the packet's sequence number was zero and no flags were set.	LOG_NOTICE
JSERVICES_TCP_SEQNUM_ZERO_FLAGS_SET	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The packet was discarded because the packet's sequence number was zero and one or more of the FIN, PSH, and URG flags were set.	LOG_NOTICE
JSERVICES_UDP_HEADER_LEN_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The UDP packet was discarded because the length field in the packet header was shorter than the minimum 8 bytes required for an UDP packet.	LOG_NOTICE
JSERVICES_UDP_PORT_ZERO	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -&gt; destination-address:destination-port, event-description</i>	The UDP packet was discarded as the source or destination port specified in the packet was zero.	LOG_NOTICE

## System Management

- **Change to process health monitor process (MX Series)**—Starting in Junos OS Release 15.1F5, the process health monitor process (pmond) is enabled by default on the Routing Engines of MX Series routers, even if no service interfaces are configured. To disable the pmond process, include the **disable** statement at the **[edit system processes process-monitor]** hierarchy level.

## Platform and Infrastructure

- **Egress Multicast Replication**—Starting with Junos OS Release 16.1, you can enable egress multicast replication to optimize multicast traffic in a Junos Fusion. In egress multicast replication, multicast traffic is replicated on satellite devices, rather than on the aggregation device. If you have a large number of multicast receivers or high multicast bandwidth traffic, enabling egress multicast replication reduces the traffic on cascade port interfaces and reduces the load on the aggregation device. This can reduce the latency and jitter in packet delivery, decrease the number of problems associated with oversubscription, and prevent a traffic storm caused by flooding of unknown unicast packets to all interfaces.

This feature is disabled by default. To enable egress multicast replication, use the **local-replication** statement in the **the [edit forwarding-options satellite]** hierarchy level. When you enable this feature, local replication is enabled on all satellite devices that are connected to the aggregation device. You cannot enable local replication for just a few selected satellite devices, specific bridge domains, or specific route prefixes.

Egress multicast replication does not take effect with the following features (Junos Fusion replicates multicast traffic on the aggregation device and other multicast traffic will continue to be replicated on satellite devices):

- Multicast support on pure layer 3 extended ports
- MLD snooping on an IPv6 network

Egress multicast replication is incompatible with the following features (the feature will not work together with egress multicast replication and you must choose either to enable egress multicast replication or to use the feature):

- VLAN tag manipulations, such as VLAN tag translations, VLAN tag stacking, and VLAN per port policies. This can result in dropped packets caused by unexpected VLAN tags.
- Multicast support for the extended ports on the edge side of Pseudowire connections in VPLS networks.
- Multicast support for the extended ports on the edge side of EVPNs.
- Multicast VPN deployments.
- MPLS/BGP VPN deployments.
- Features that perform egress actions on individual extended ports, such as egress local-port mirroring.

Use the following new operational commands to display information related to this feature:

- **show bridge flood next-hops satellite**
- **show bridge flood next-hops satellite nexthop-id *nexthop-identifier***
- **show bridge flood satellite**
- **show bridge flood satellite bridge-domain-name *domain-name***
- **show bridge satellite device**
- **show multicast ecid-mapping satellite**
- **show multicast next-hops satellite**
- **show multicast snooping next-hops satellite nexthop-id *nexthop-identifier***
- **show multicast snooping route satellite**
- **show multicast snooping route satellite bridge-domain-name *domain-name***
- **show multicast snooping route satellite group *group-id***

- **show multicast statistics satellite**
- **show multicast summary satellite**
- The length of TACACS messages allowed on JUNOS devices has been increased from 8150 to 65535 bytes. [PR1147015](#)

### Virtual Chassis

- **SNMP MIB walk on MX Series Virtual Chassis**—Starting with Junos OS Release 15.1F5, **snmp mib walk** operations no longer return invalid PCMCIA card information for Routing Engines on MX Series Virtual Chassis.

### VPNs

- **Clear all Internet key exchange (IKE), traffic encryption key (TEK), key encryption key (KEK), and security associations (SAs) for group VPN (MX Series)**—The **clear security group-vpn member group** CLI command has been introduced in the Release 15.1F3 of Junos OS for MX Series routers to clear all Internet key exchange (IKE), traffic encryption key (TEK), key encryption, and key (KEK) security associations (SAs) for a group VPN.

```
user@host> clear security group-vpn member group
```

- See Also**
- *New and Changed Features*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Known Issues*
  - *Resolved Issues*
  - *Documentation Updates*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Product Compatibility*

## Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1F7 for the MX Series.

- [Hardware on page 52](#)
- [General Routing on page 52](#)
- [Interfaces and Chassis on page 54](#)
- [MPLS on page 55](#)
- [OpenFlow on page 55](#)
- [Services Applications on page 55](#)

- [Subscriber Management and Services \(MX Series\) on page 55](#)
- [User Interface and Configuration on page 56](#)

## Hardware

---

- **Support for MIC-3D-8OC3-2OC12-ATM Revision 22 and later**—MIC-3D-8OC3-2OC12-ATM Revision 22 is supported only by the following Junos OS releases:
  - Junos OS Release 12.3—12.3R9 and later
  - Junos OS Release 13.3—13.3R6 and later
  - Junos OS Release 14.1—14.1R4 and later
  - Junos OS Release 14.2—14.2R3 and later
  - Junos OS Release 15.1 and later

You must upgrade to a supported Junos OS release to use MIC-3D-8OC3-2OC12-ATM Revision 22 and later.

- **The options alarm low-light-alarm and warning low-light-warning might not work (MX Series)**—The **alarm low-light-alarm** and **warning low-light-warning** options at the **[edit interfaces interface-name optics-options]** hierarchy level might not work for the 10-Gigabit Ethernet and 100-Gigabit Ethernet interfaces of MPC3, MPC4, MPC5, MPC6, MPC7E, MPC8E, and MPC9E on MX Series 3D Universal Edge Routers. These options might not work on MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q if they are installed with Junos Continuity software.

This is a known behavior and has no impact on the performance of these line cards.

## General Routing

---

- The following **request** commands are not available for the Routing Engine RE-S-X6-64G on the MX240, MX480, MX960, MX2010, and MX2020:
  - **request system halt**
  - **request system partition**
  - **request system power off**
  - **request system power on**

The scope of functionality of the following commands is limited to Junos OS guest level:

- **request system reboot**
- **request system snapshot**
- **request system software add**
- **request system zeroize**

You can use the following equivalent **request vmhost** commands to achieve the functionality:

- request vmhost cleanup
- request vmhost file-copy
- request vmhost halt
- request vmhost hard-disk-test
- request vmhost power-off
- request vmhost power-on
- request vmhost reboot
- request vmhost snapshot
- request vmhost software abort
- request vmhost software add
- request vmhost software in-service-upgrade
- request vmhost software rollback
- request vmhost zeroize

The output of the following **show** commands are modified for the Routing Engine RE-S-X6-64G:

- show chassis environment routing-engine
- show chassis hardware
- show chassis hardware extensive
- show chassis routing-engine
- show system software

The following new **show** commands are introduced for the Routing Engine RE-S-X6-64G:

- show vmhost bridge
- show vmhost crash
- show vmhost hardware
- show vmhost information
- show vmhost logs
- show vmhost netstat
- show vmhost processes
- show vmhost resource-usage
- show vmhost snapshot
- show vmhost status

- **show vmhost uptime**
- **show vmhost version**

The following new configuration statements are introduced for the Routing Engine RE-S-X6-64G:

- **edit system processes app-engine-virtual-machine-management-service**
- **edit vmhost**
- During deletion and restoration of scaled configurations on MX240, MX480, MX960, MX2010, and MX2020, error messages related to next hops are displayed.
- During graceful Routing Engine switchover (GRES) on MX240, MX480, MX960, MX2010, and MX2020, the 802.1x daemon crashes if the number of logical interfaces configured is equal to 64,000 or more. After the crash, the daemon restarts and resumes normal operations.
- The configuration of the smartd process, which monitors the status of the disk on the host OS of MX240, MX480, MX960, MX2010, and MX2020, is not deleted completely even after you delete the configuration. When you configure the smart check feature, smartd continues to use parameters that were configured previously. Therefore, while enabling smart check, remember to configure the threshold values for smartd instead of retaining the default values that were previously configured.
- FIFO handles of SSD-monitoring smartd are not cleared on the host OS after multiple commits or checks. Smartd stops working when the FIFO limit reaches a maximum. Therefore, we recommend that you do not change smartd configurations too often and perform SSD smart checks after long intervals of time. When the FIFO limit reaches a maximum, reboot the host OS.
- In a dual Routing Engine system, while Junos OS has just started booting in the master Routing Engine, the backup Routing Engine might be powered-off if it is removed and reinserted.

As a workaround, plug in the backup Routing Engine after master Routing Engine is running stable and all the FPCs are in operational state. If the other Routing Engine gets powered-off accidentally, issue the commands **request chassis cb slot number power off** and **request chassis cb slot number power on** to turn the power on the Routing Engine. The *slot number* signifies the Routing Engine that has to be powered-on.

- The date and time zones are synchronized from the admin guest Junos OS to host OS on the MX240, MX480, MX960, MX2010, and MX2020 and use same time zones. Therefore, there is no difference in the timestamp in system log files of Junos OS and the host OS.

---

## Interfaces and Chassis

- Starting in Junos OS Release 15.1F4, when interfaces are disabled on MPC7, the output of the **show interfaces diagnostic optics** command displays the following information under lane characteristics:

Tx laser disabled alarm : Off/On

## MPLS

- The configuration **flow-label-transmit** and **flow-label-receive** statements are not supported in OAM CFM session over L2Circuit.

## OpenFlow

- On MX Series routers running OpenFlow v1.3.1 or later, a group in which the same output port is specified for multiple buckets is not supported.

## Services Applications

- The MS-MPC cannot support the route scale of other MPCs because the MS-MPC does not have as much memory as other MPCs.

## Subscriber Management and Services (MX Series)

- By default, Link Aggregation Control Protocol link protection is revertive. This means that after the current link becomes active, the router switches to a higher-priority link if one becomes operational or is added to the aggregated Ethernet bundle. In a highly scaled configuration over aggregated Ethernet, we recommend that you prevent the router from performing such a switch by including the **non-revertive** statement at the **[edit chassis aggregated-devices ethernet lacp link-protection]** hierarchy level. Failure to do so may result in some traffic loss if a MIC on which a member interface is located reboots. Using the **non-revertive** statement for this purpose is not effective if both the primary and secondary interfaces are on the MIC that reboots.
- **Dynamic firewall filter match conditions for enhanced subscriber management (MX Series)**—Enhanced subscriber management does not support dynamic firewall filter match conditions that consist of an interface identifier and an \* (asterisk) wildcard character, such as **interface pp0.\*** or **interface demux0.\***. If you use interface specifications with wildcards as match conditions, the match results do not include dynamic subscriber interfaces created with enhanced subscriber management. However, interface specifications that use a wildcard character continue to be supported for statically configured interfaces.

In earlier Junos OS releases, match conditions consisting of an interface identifier and an \* (asterisk) wildcard character are supported for both dynamically configured and statically configured interfaces.

- **Support for multicast group membership in Enhanced Subscriber Manager (MX Series)**—In Junos OS Release 15.1F5, enhanced subscriber management does not support the use of dynamic profiles for the static configuration of multicast group membership for subscribers. Instead, subscribers must send an IGMP JOIN message to receive the multicast stream. More specifically, the following command is not supported in this release:

```
set dynamic-profiles client profile protocols igmp interface $junos-interface-name static group 224.117.71.1
```

- **Dynamic Provisioning in Layer 2 Wholesaling (MX Series)**—Starting with release 15.1R3, Junos OS does not support dynamic VLAN mapping into VPLS instances (you

can however still configure static VLAN interface mapping to VPLS instances). By extension, this means that dynamic provisioning for Layer 2 wholesaling is also not supported in the current release. Note that dynamic provisioning in layer 2 wholesaling is not being dropped from Junos; support is slated to return in a future release.

The following sample shows the exact commands that are not currently available (encapsulation VLAN-VPLS, and family VPLS, under the dynamic interfaces hierarchy):

```
interfaces {
    "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
            encapsulation vlan-vpls
                vlan-id "$junos-vlan-id";
            family vpls;
        }
    }
}
```

---

### User Interface and Configuration

- **Modification to configurable link degrade threshold values (MX Series)**—Starting with Junos OS Release 15.1F7, the values of the user configurable link degrade thresholds, have to be configured as per the following guidelines:
  - **set threshold value** must be greater than **warning set threshold value**
  - **set threshold value** must be greater than **clear threshold value**
  - **warning set threshold value** must be greater than **warning clear threshold value**

If the threshold values are not configured as per these guidelines, the configuration fails and a Commit Error message is displayed.

- See Also**
- *New and Changed Features*
  - *Changes in Behavior and Syntax*
  - *Known Issues*
  - *Resolved Issues*
  - *Documentation Updates*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Product Compatibility*

### Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1F7 for the MX 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 \(CoS\) on page 57](#)
- [Forwarding and Sampling on page 57](#)
- [General Routing on page 58](#)
- [High Availability \(HA\) and Resiliency on page 63](#)
- [Infrastructure on page 63](#)
- [Interfaces and Chassis on page 63](#)
- [Layer 2 Features on page 64](#)
- [Layer 2 Ethernet Services on page 64](#)
- [Multiprotocol Label Switching \(MPLS\) on page 65](#)
- [Platform and Infrastructure on page 66](#)
- [Routing Protocols on page 67](#)
- [Services Applications on page 69](#)
- [Software Installation and Upgrade on page 69](#)
- [Subscriber Access Management on page 69](#)
- [User Interface and Configuration on page 70](#)
- [VPNs on page 70](#)

---

### Class of Service (CoS)

- When the "chained-composite-next-hop" is enabled for Layer 3 VPN routes, MPLS CoS rewrite rules attached to the core-facing interface for "protocol mpls-inet-both-non-vpn" are applied not only to non-VPN traffic (which is the correct behavior) but also to Layer 3 VPN traffic. That is, both MPLS and IP headers in Layer 3 VPN traffic receive CoS rewrite. [PR1062648](#)
- In rare cases, cos related queue stats polling with multiple OID packing or multiple snmp client polling on same interface simultaneously can cause cosd coredump and restart. cosd restart does not impact any cos services. [PR1199687](#)
- Following error log message might be seen with Hierarchical CoS and strict-high scheduling configured. Dec 27 11:08:02.293 mand-re0 fpc1  
cos\_check\_temporal\_buffer\_status: IFD ge-1/2/1 IFL 358: Delay buffer computation incorrect.^M If hierarchical scheduler is configured for an IFD and if guaranteed rate is not set for an IFL under this IFD, then the temporal buffer configured The display of error message is valid when guaranteed rate is '0', but it is not valid when 'guaranteed rate' is disabled. [PR1238719](#)

---

### Forwarding and Sampling

- The "default-arp-policer" is applied to every relevant IFL to rate limit the ARP traffic. You can disable the "default-arp-policer" by running the above hidden command "set firewall disable-arp-policer". Note that improper application leads to the Routing Engine over loaded with a bulk of ARP traffic leading to a typical DOS scenario. The

issue was that even after disabling the "default-arp-policer", it still affected IFL in some scenario such as after DUT reboot or when a new IFL is created. The issue is fixed in this PR so that wherever "set firewall disable-arp-policer" is configured, in all scenarios "default-arp-policer" will not get applied to IFL. [PR1198107](#)

- After upgrading to 15.1F5 or 15.1F6 DFWD is at 100 percent on backup Routing Engine. This is Fixed in 15.1F7 and later. [PR1219562](#)
- Bandwidth-percent policer does not work on ps interface, which will result in commit error. [PR1225977](#)

### General Routing

---

- When ps interface is configured using as anchor interface a logical tunnel (lt) interface without explicit tunnel-bandwidth configuration (under 'chassis fpc <fpc-number> pic <pic-number> tunnel-services' configuration hierarchy), the ps interface is created only in kernel, but not on Packet Forwarding Engine. In order to have ps interface in Packet Forwarding Engine, an explicit tunnel-bandwidth configuration is required. PR 1042737 removes this restriction, and a ps interface may be anchored to an lt interface without explicit tunnel-bandwidth configured. [PR1042737](#)
- DPD/IKEv2 informational messages are dropped at the peer. Adding vendor ID in the INFORMATIONAL message is causing peer to drop such packets. [PR1066336](#)
- The configuration support for enabling ingress and egress layer2-overhead is available in dynamic-profile but the functionality is not supported in 15.1R3 and 15.1R4. For example, set interfaces ge-4/2/9 unit 0 account-layer2-overhead ingress 30 set interfaces ge-4/2/9 unit 0 account-layer2-overhead egress 30 With the above configuration, the number of layer2-overhead bytes (30) are not added to the input bytes in traffic statistics. [PR1096323](#)
- For Junos 13.3R1 and later, the DPC card might experience a performance degradation when it's transferring bidirectional short packets (64B) in inline rate. [PR1098357](#)
- In a multicast virtual private network (MVPN) scenario during route churn, the rpd process might crash due to inconsistency multicast next-hop between rpd and kernel. [PR1138366](#)
- In sampling feature, certain scenarios force handling of the sampled packet at the interrupt context, which may have chance to corrupt the BMEB packet context, and lead to BMEB FDB corruption. [PR1156464](#)
- A previous enhancement to strengthen the VC-Heartbeat message exchange resulted rejecting messages at the crucial time of determining the health of the other VC member when all adjacency links fail. Validation of messages has been adjusted to remain strong when the VC is connected, but relaxed during the split conditions to prevent rejecting valid messages. [PR1157383](#)
- On MX Series routers with MS-MICs and MS-MPCs with the syslog statement included at the [edit services cos rule rule-name term term-name then] hierarchy level, a system log message is not generated when a CoS rule term is matched, in contrast to the expected behavior in which system log messages are generated when a NAT rule term is matched. [PR1159231](#)

- This is an intermittent issue. Assuming that AE is configured with the bypass-queuing-chip configuration statement. Now followup configuration changes are such that removing child link(s) from AE bundle, configuring per-unit-scheduler on the removed child link(s) in a single commit causes intermittent issues with per-unit-scheduler configuration updates to cosd and the Packet Forwarding Engine. Hence, dedicated scheduler nodes might not be created for all units or IFLs. [PR1162006](#)
- MQCHIP reports continuous "FI Cell underflow at the state stage" message and continuous fabric drops on ADPC ICHIP Packet Forwarding Engines after ISSU on MX with ADPC. [PR1163776](#)
- Memory leak in kernel caused by dot1xd that leads to system crash [PR1163782](#)
- On MS-MIC, starting from 15.1R3 onwards, the JFLOW/Sampling scaling is coming down to 12.5 million active flows. [PR1163976](#)
- When upgrading Junos software on RE1, if at the time, RE1 is the "master RE", both REs will be in "backup" state. Resulting in losing remote connectivity, and all interfaces. Only "console" access will be available at this time. [PR1172729](#)
- MTU discovery may not be working due to lack of VRF info on egress card for BBE Subscriber traffic. [PR1177381](#)
- This is a display issue and doesn't affect functionality of the power, fixing has been added to commands 'show chassis power' and 'show chassis environment pem', when one of the DC PEM circuit breaker tripped. [PR1177536](#)
- If "router-advertisement" protocol is configured in client ppp profile, unsolicited RA might be sent before the IPv6CP Configuration ACK is received. [PR1179066](#)
- With basic-nat44 or twice-basic-nat-44 NAT translation type on ms-mic/ms-mpc, it is not possible to use the same pool against different base address ranges. In other words, it is not possible, within the same service-set, to use two different terms with a different 'from' address criteria on the same pool. [PR1183217](#)
- AMS redundant interfaces not listed under possible-completions of operational commands. [PR1185710](#)
- static ifl (static dhcp i.e dhcp not by demux or ipdemux dynamic profile) cannot be added to a dynamic-profile created IFLSET. subscriber login attempt on such a config where static ifl is added over a dynamic-profile iflset will be rejected. static ifl for example is dhcp over vlan interface without any demux or ipdemux. [PR1185956](#)
- CGNAT NAT mappings and ports are not cleared after SIP session timeouts for the SIP spoofed traffic and SIP scaled traffic, and very few ports not released for the scaled traffic of 10-20K SIP sessions. [PR1187965](#)
- JUNOS might improperly bind PFE ukernel application sockets after ISSU due to a bug in IP->TNP fallback logic. Because of that bug, threads running on the ukernel that relay on UDP sockets can experience connectivity issues with host, which in turn can lead to various problems. For instance, sntp (simple network time protocol) client might fail to synchronize time, which in turn might lead to other problems such as failure in adjacency formation for HMAC authenticated protocols. [PR1188087](#)

- The rdd (a daemon used by MS-DPC/MS-MPC/AE) might crash after performing certain operations on dual RE MX platform with AE interface configured and non-Trio FPC installed. [PR1188832](#)
- Due to a bug in schema with Junos Version 14.1Rx and 15.1rx Releases, Admins will not be able push mpls configurations to Devices which includes loose strict tags [PR1193599](#)
- Line card is crashed while doing internal testing. Line card is busy looping/not yielding to other threads so chassisD sent an NMI and crashed the line card. This is observed only once and unable to reproduce. [PR1194692](#)
- If a fragmented ICMP request from subscribers is sent to a device, the device only responses ICMP request for the first packet, which causes PING failing. And if PING initiates from a device to subscribers with size greater than the negotiated MRU, the device can not fragment the packet, which causes PING still failing. [PR1195031](#)
- During ISSU (unified in-service software upgrade) it might be noticed that the below log messages are displayed: SFP: pointer Null, sfp\_set\_present This might trigger a flap in the interfaces on MX routers while upgrading using ISSU. [PR1200045](#)
- When performing unified in-service software upgrade (ISSU) on MX series routers, the MPC might crash during the field-replaceable unit (FRU) upgrade process. [PR1200690](#)
- The routing table will not be updated if some of the unnumbered interface goes down and some unnumbered interface is still active when there are multiple unnumbered interfaces configured under OSPF. [PR1202795](#)
- When dynamic firewall filter is configured to match packet in prefix/mask format, the firewall filter might not be correctly programmed on PFE. [PR1204291](#)
- Interface statistics shows Last cleared : Never even after clearing the statistics even though the stats are cleared using CLI. The fix will be resolved in upcoming Junos releases 15.1R5 and 15.1R6 [PR1206407](#)
- In a fully loaded MX2008 chassis, the overall time taken for all 8 SFB2s to come online from offline state is around 15mins. Which is much more compared to MX2010 SFBs which takes around 6mins to come online. [PR1207581](#)
- VC link "last flapped" timestamp is reset to "Never" on the new backup RE after MX VC global GRES switchover. [PR1208294](#)
- When trying to scale total numbers of subscribers on a chassis beyond 375K with 4 MPC5E-Q cards in an MX 480/960 chassis, the clients may get rejected due to memory threshold being exceeded. The resource monitoring output shows incorrect value for expansion memory. The system will still allow 128K subscribers to be scaled on a single MPC5 cards. [PR1210122](#)
- AMS basically aimed at, one active service PIC failure at a time in N:1 (N >1) redundancy. In the event more than one active service PIC fails, in case of NAT service, respective NAT pool belonging to the second failed service PIC, would not be used till AMS reaches a state where there is only one or no service PIC failed. [PR1210187](#)
- On setup with IRB config, when actions which result into underlying AE interface of IRB going down, are done, the backup RE may experience 'panic' and hence reboot. The panic will be due to not being able to allocate nexthop index that master RE has asked.

Since the panic and reboot happens on backup routing/forwarding/any other functionality will not be affected. Some examples of trigger - continuous child link flaps of AE or back-to-back commits of different IRB configs etc. [PR1211900](#)

- The MS-MPC/MS-MIC service cards might encounter a core when using certain ALGs or the EIM (Endpoint-independent mapping ) / EIF (Endpoint independent filtering) feature due to a bad mapping in memory. [PR1213161](#)
- Inline Jflow service will not work after ISSU on MPC5E and above type line cards. [PR1214842](#)
- MX-VC: All VCP interface experiences tail-dropped as result of configuration conflict. It is a good idea to reference documentation and customize the COS associated with VCP interfaces. In this scenario customer has configured a corresponding xe-n/n/n interface with just a description to denote that port is dedicated to VCP. Problem is the resource calculation is impacted and reports smaller queue-depth maximum values when both network interface xe-n/n/n and vcp-n/n/n are defined. Issue is more likely to occur with dynamic modification add/delete of vcp interfaces with a corresponding network interface xe-n/n/n configured. > show interfaces queue vcp-5/3/0 | match max Maximum : 32768 Maximum : 32768 Maximum : 32768 Maximum : 32768 [PR1215108](#)
- On 15.1R3 onwards MX trinity platform release, if DHCPv4 or DHCPv6 subscriber is configured and the subscriber joins more than 29 multicast groups, the line card might crash. [PR1215729](#)
- Incorrect source MAC used for PPPoE after underlying AE is changed [PR1215870](#)
- In the case of multi-homed (MH) PEs with EVPN, there is a possibility of an rpd core during MAC move between MH PEs. [PR1216144](#)
- The bbe-smgd core occurred in bbe\_autoconf\_if\_l2\_input when DHCP client generates ARP. [PR1220193](#)
- There is no ISSU from 15.1 and older releases to 16.2R1. [PR1222540](#)
- In some scenarios Service Name Table not found: name" error is displayed after commit and service name configuration is not applied. [PR1222551](#)
- During CoA request there are no changes on schedulers. Requests are received successfully, but no changes from CoS side. [PR1222553](#)
- The problem of tunnel stream getting mis configured for LT interfaces is due to internal programming and the same has been corrected to evaluate multiple lt interfaces for FPC and PIC slot combination. [PR1223087](#)
- In PPPoE subscriber scenario, after demux underlying interface AEx is changed to AEy, the source MAC used for PPPoE handshake is still the old AEx interface's MAC. This causes PPPoE clients to fail as the PADR packets from the client are dropped due to the MAC address mismatch. [PR1224190](#)
- CPU utilization is calculated per MS-MIC. Under each MS-MIC, we can have ms-\* and mams-\* interfaces. Total CPU utilization is shown under jnxSpSvcSetIfCpuUtil of ms-\*, even if its configured as mams-\* interface. [PR1227004](#)

- Due to a bug in JUNOS code, the subscriber's traffic volume accounting stats remains unchanged post-ISSU on MXVC platform until exceeds the pre-ISSU value. This is a day-1 issue seen on MXVC after JUNOS 14.1. [PR1230524](#)
- Unsuccessful DCE-RPC ALG sessions results in stale gates and lead to MS-MPC/MS-MIC restart when the gates clean-up occurs after timeout. [PR1230868](#)
- In an idle MX2008 Chassis, chassisd CPU utilization oscillates between 10% to 20%. There will be no functional impact. [PR1231333](#)
- A wrong PE is being attached to an ESI when the router receives two copies of the same AD/ESI route (e.g. one through eBGP and another one received from an iBGP neighbor). This will causes partial traffic blackhaule and stale MAC entries. You can confirm the issue by checking the members of the ESI: labroot@MX2> show evpn instance extensive ... Number of ethernet segments: 5 ESI:  
00:13:78:00:00:00:00:00:01 Status: Resolved Number of remote PEs connected:  
3 Remote PE MAC label Aliasing label Mode 87.233.39.102 0 0 all-active 87.233.39.1  
200 0 all-active <<<< this PE is not part of the ESI 87.233.39.101 200 0 all-active  
[PR1231402](#)
- ICMP Identifier not translated back to expected value during traceroute for TTL exceeded packets on NAT using Multiservice MPC. This occurs for ICMP ID >255 and causes all hops (except 1st and last) appearing as "\*". [PR1231868](#)
- When set port-mirror to MX router, LSP-ping might fail and IP packets with options will not get mirrored due to unexpected echo reply from DUT:  
<----- echo request -----> echo reply  
[R1]-----[DUT]-----[R2] A | -----> echo reply (unexpected behavior)  
| mirror [PR1234006](#)
- MX MPC7 and above might receive noise on the FPC console port, and interprets it as valid signals. This might cause login fails on the console port, core-dumps or even reloads. This PR covers MX cards, PR/1224820 covers PTX FPC3. [PR1234712](#)
- VLNS(VBNG) - Commit generated a "warning: requires 'l2tp-inline-lns' license" but valid license is installed [PR1235697](#)
- When per-packet load sharing is enabled under Aggregated Ethernet interface, egress traffic over the Aggregated Ethernet interface might be dropped unexpectedly. [PR1235866](#)
- DNS server IP addresses are not present in the output of 'show subscribers extensive' for DHCP subscribers in case DNS configuration is provided from the access-profile or pool. In case of such data is provided from RADIUS, the output is correct. The issue is cosmetic: DNS addresses are provided to subscribers. [PR1237525](#)
- trace route won't resolve VRF loopback address where SI and Pseudo interface exist [PR1240221](#)
- When CLK0 and/or CLK1 port of MX2008 is configured as BITS clock source/s, BITS LED glows GREEN when external BITS clock is not connected to the port. When external BITS clock is connected and then removed BITS LED glows RED. [PR1222041](#)
- JCS port testing functionality will not be supported on the two optics ports present on the Front faceplate of the CB-RE. [PR1208638](#)

- ISSU (In Service Software Upgrade) will not be supported. [PR1213193](#)
- On yankout of MPC, Link Error column in show chassis fabric summary extended output shows YES for all fabric planes. Whereas, on offline of MPC using cli command, output shows correctly. [PR1214611](#)

### High Availability (HA) and Resiliency

- When B2B switch over is done on TXP with 15.1F6 image, on first switch over the system performs as expected. However packet loss may be seen on doing switch over for the second time. Here in second switch over , 0.21% packet loss happening. [PR1172546](#)

### Infrastructure

- LACP daemon runs only on the current master RE and delegates sending and receiving PDUs to PPMAN. During GRES, PPMAN holds these xmit and rcv entries until the LACP daemon comes up on the new master and re-programs these xmit and rcv entries to PPMAN. The maximum interval for which PPMAN holds these entries is 15s by default. In this case, LACP daemon is unable complete start-up events and re-program PPMAN within 15s, and hence PPMAN stops sending LACP PDUs post 15s. This leads to timeout on the remote end and the AE interfaces flap. [PR1202622](#)
- Unable to execute 'show log user' cmd because freebsd 'last' utility is not exist in FreeBSD 10.x-based Junos OS package [PR1221581](#)

### Interfaces and Chassis

- AE interfaces may go down during a MX-VC ISSU procedure in a scaled system, leading to traffic and subscriber loss. [PR1191909](#)
- When VRRP is configured on IRB interface with scaling configuration (300k lines), in corner case, handles might not be released appropriately after their use is over. As a result of that, memory leak on vrrpd might be seen after configuration commit. [PR1208038](#)
- During L2TP session establishment on MX LAC, if CPE attempts to negotiate MRU higher than 1492 bytes, spurious MRU of 1492 bytes is included into the Last Received ConfReq AVP in ICCN packet. [PR1215062](#)
- In ppp subscriber scenario, if the jpppd process receives a reply message attribute from the radius or tacplus server with a character of %, it might cause the jpppd process to crash and cause the ppp user to be offline [PR1216169](#)
- Dcd can not start after router reboot due to non-existing IFL referenced in 'demux-options underlying-interface' [PR1216811](#)
- The configuration change where for a static vlan demux interface the underlying physical interface is changed to a one with lower bandwidth (e.g. from xe to ge) can fail with the following error: "error: Bandwidth on IFL demux0.7000 cannot be greater than that of its IFD". For example: user@router# show | compare [edit interfaces demux0 unit 7000 demux-options] - underlying-interface xe-0/1/0; + underlying-interface ge-0/3/9; lab@jovian-re0# commit re0: error: Bandwidth on IFL demux0.7000 cannot be greater

than that of its IFD error: DCD Configuration check FAILED. error: configuration check-out failed [PR1232598](#)

- In case if there is a iflset configuration present then the following issue might be seen with the windsurf card interfaces - After ISSU from non-FreeBSD 10.x-based Junos OS to 15.1F throttle, interfaces of windsurf card stay down. when the card is restarted, it goes to ready state. - After ISSU from non FreeBSD 10.x-based Junos OS to 16.1 throttle, windsurf card interfaces stay down but neo card goes to ready state. [PR1242627](#)
- In case more than one IFL (logical interface) are configured under same IFD (physical interface), and VRRP is configured on one IFL without vlan and the lower unit number IFL has vlan config present, then vrrpd wrongly carries vlan info from the lower unit number IFL to this IFL's configuration. And so VRRP might be stuck in (state: unknown, VR State: bringup). This might happen in the scenario of VRRP configured on physical interface with flexible-vlan-tagging or lt interface without flexible-vlan-tagging. [PR1247050](#)

---

## Layer 2 Features

- When "input-vlan-map" with "push" operation is enabled for dual-tagged interfaces in "enhanced-ip" mode, there is a probability that the broadcast, unknown unicast, and multicast (BUM) traffic may be blackholed on some of the child interfaces of the egress Aggregated Ethernet (AE) interfaces or on some of the equal-cost multi-path (ECMP) core-links. [PR1078617](#)

---

## Layer 2 Ethernet Services

- IPv4 and IPv6 long Virtual Router Redundancy Protocol (VRRP) convergence delay and unexpected packet loss might happen when MAC move for the IRB interface occurs (e.g. when flapping the Layer 2 interface which is the under-interface of IRB on master VRRP). [PR1116757](#)
- In DHCP environment, if interface is deleted and recreated in single commit, the duplicate DHCP subscriber is not getting bound. [PR1188026](#)
- If a client sends a DHCP Request packet, and Option 55 includes PAD option (0), a DHCP ACK will not be sent back to the client. [PR1201413](#)
- When MSTP is configured under routing-instance, both the primary and standby VPLS pseudowires are stuck in ST state due to a bug in the software. That has been fixed and now the PW status is set correctly. [PR1206106](#)
- In DHCP relay environment, when delay-authentication and proxy-mode are configured at same time, jdhcpd may crash due to NULL session ID. [PR1219958](#)
- MX is not including Delegated-IPv6-Prefix in accounting interim. [PR1231665](#)
- This issue can be seen if CPE is initiating DHCPv6-Solicit with IA\_NA, IA-PD and Rapid-Commit Option but MX will send the DHCv6 Advertise with Rapid commit flag even though Rapid-commit knob is not enabled on MX. [PR1235578](#)



## Multiprotocol Label Switching (MPLS)

- When graceful Routing Engine switchover (GRES) is done between the master and backup Routing Engines of different memory capabilities (such that one has only enough memory to run routing protocol process (rpd) in 32-bit mode while the other is capable of 64-bit mode, which could be caused by using Junos OS Release 13.3 onwards with the configuration statement "auto-64-bit" configured, or, using Junos OS Release 15.1 onwards even without the configuration statement), rpd might crash on the new master Routing Engine. As a workaround, this issue could be avoided by the CLI command "set system processes routing force-32-bit". [PR1141728](#)
- From Junos OS Release 16.1R1 with ldp egress protection in stub-alias mode, traffic loss occurs when the interface between the protector egress node and the primary egress node goes down. [PR1190983](#)
- With a high degree of aggregation and a large number of next hops for the same route, LDP may spend too much CPU updating routes due to topology changes. This may result in scheduler slip, ldp session timing out and long LSP convergence. [PR1192950](#)
- In impacted Junos releases ldp will import metric for all isis routes which have tags without knob track-igp-metric. Versions 14.1R3, 14.2R1 and later are impacted with this issue. For example below route has tag and ldp metric is same as IGP MX> show route 20.20.20.20/32 inet.0: 17696 destinations, 17696 routes (17695 active, 0 holddown, 1 hidden) + = Active Route, - = Last Active, \* = Both 20.20.20.20/32 \*[IS-IS/18] 3d 03:36:14, metric 11, tag 80 to 10.10.1.1 via xe-0/0/2.0 to 10.11.1.1 via xe-0/0/3.0 > to 10.13.1.1 via xe-2/1/2.0 to 10.14.1.1 via xe-2/1/3.0 inet.3: 13418 destinations, 13418 routes (13418 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, \* = Both 20.20.20.20/32 \*[LDP/9] 00:58:01, metric 11, tag 80 to 10.10.1.1 via xe-0/0/2.0, Push 533376 to 10.11.1.1 via xe-0/0/3.0, Push 533376 to 10.13.1.1 via xe-2/1/2.0, Push 533376 > to 10.14.1.1 via xe-2/1/3.0, Push 533376 Below route does not have tag and has default Ldp metric MX> show route 10.10.10.10/32 inet.0: 17695 destinations, 17695 routes (17694 active, 0 holddown, 1 hidden) + = Active Route, - = Last Active, \* = Both 10.10.10.10/32 \*[IS-IS/18] 3d 03:35:23, metric 22 to 10.1.1.1 via xe-0/0/0.0 to 10.10.1.1 via xe-0/0/2.0 to 10.11.1.1 via xe-0/0/3.0 to 10.12.1.1 via xe-2/0/0.0 > to 10.13.1.1 via xe-2/1/2.0 to 10.14.1.1 via xe-2/1/3.0 inet.3: 13417 destinations, 13417 routes (13417 active, 0 holddown, 0 hidden) + = Active Route, - = Last Active, \* = Both 10.10.10.10/32 \*[LDP/9] 00:57:10, metric 1 to 10.1.1.1 via xe-0/0/0.0, Push 593760 to 10.10.1.1 via xe-0/0/2.0, Push 556161 to 10.11.1.1 via xe-0/0/3.0, Push 556161 to 10.12.1.1 via xe-2/0/0.0, Push 593760 to 10.13.1.1 via xe-2/1/2.0, Push 556161 > to 10.14.1.1 via xe-2/1/3.0, Push 556161 [PR1225592](#)
- Routing Protocol process (RPD) may stop running unexpectedly if a static MPLS LSP is moved from one routing instance to another routing instance in one single configuration change with one single commit. RPD will need a manual restart with "restart routing" [PR1238698](#)

## Platform and Infrastructure

---

- When next-ip is defined as the action and there is no ARP for the IP address specified under next-ip, the traffic is not forwarded. A manual ping needs to be initiated for things to work. [PR864861](#)
- In the dual Routing Engine scenario, the backup Routing Engine does not sync up the configuration change while deleting an inactivated interface from the master. So after the operation, the inactivated interface still exists on the backup Routing Engine. [PR991081](#)
- In configurations with IRB interfaces, during times of interface deletion, such as an FPC reboot, the Packet Forwarding Engine may log errors stating "nh\_ucast\_change:291Referenced l2ifl not found". This condition should be transient, with the system re-converging on the expected state. [PR1054798](#)
- SNMP queries to retrieve jnxRpmResSumPercentLost will return the RPM/TWAMP probe loss percentage as an integer value whereas the precise value (including decimal points) can be retrieved through the CLI by using the following commands: show services rpm probe-results show services rpm twamp client probe-results [PR1104897](#)
- When "persist-groups-inheritance" is configured, mgd process is not setting CHANGED bit in config DB for [policy-options prefix-list <> apply-path] correctly. So rpd process thinks "apply-path" hasn't changed and doesn't read this config-path again. And apply-path function is broken. [PR1173443](#)
- HTTP connections are not successful when AE + ECMP between PE and P with servicing(sfw+nat+cos) being done on MS-DPC. This issue is still under investigation.... Investigation Update (9/2/2016): LAG Enhanced is not compatible with chained-composite next-hops on the MS-DPC. Either LAG enhanced or chained-composite nexthops can be disabled. [PR1176879](#)
- If igmp snooping is configured in a VPLS routing instance and the VPLS instance has no active physical interfaces, multicast traffic arriving from the core might be send to the Routing Engine. As a result, host queues might get congested and it might cause protocol instability. As a workaround, configure a dummy activate interface in the VPLS routing instance can avoid this issue. [PR1183382](#)
- Syslog storage in a file could abruptly stop due a race condition in handling log file rotation. The fix is available from Junos 16.1R2 onwards [PR1195239](#)
- On MX with MQCHIP linecard (Trio linecard) with traffic-control-profile, if the overhead-accounting is configured with negative values, it might not work. The shape function will be affected. [PR1195866](#)
- Unable to connect tunneled PPPoE subscribers, L2TP sessions are not getting established. The rewrite rule in dynamic-profile is not valid to attempt a ieee-802.1 VLAN rewrite on si-x/x/x in-line service interface for LTS and L2TP LNS functionality. Earlier JUNOS releases like 13.x and 14.x would ignore the configuration statement, but in release 15.1R4 the connection attempt is rejected. [PR1205883](#)
- Starting with JUNOS software version 15.1R5, 16.1R3 or later, a new mechanism is added into packet forwarding engine (PFE) to improve forwarding performance. A noticeable behavior of the mechanism is an increase in CPU utilization periodically. The results of

this mechanism is increasing in overall packet forwarding performance or better throughput. [PR1207532](#)

- On MX platforms with MPC2 NG/MPC3 NG/MPC3/MPC4/MPC5/MPC6 installed, when configuring multiple lt interfaces with HQOS on a MPC, due to a software defect, when creating internal lt tunnel stream in PFE, the tunnel bandwidth will be overridden to max bandwidth(60G for MPC2 NG/3 NG, 100G for MPC/3/4/5/6), this causes that all the 256 internal FIFO resources are only allocated for two tunnels, the allocation for other tunnels fails due to lack of resources. As a result, only two lt interfaces can stay up, other lt interfaces will go down. [PR1209065](#)
- Several files are copied between REs during 'ffp synchronize' phase of the commit (e.g. /var/etc/mobile\_aaa\_ne.id, /var/etc/mobile\_aaa\_radius.id, etc). These files are copied even if there was no corresponding change in the configuration thus unnecessarily increasing commit time. [PR1210986](#)
- If a Multicast source sends a fragmented packet (a packet which exceeds the MTU of its outgoing interface) to the router and it needs to resolve the destination route, then only the first fragment of the packet is sent when the route is resolved. [PR1212191](#)
- IPv6 traffic learned on a L2/bridge/multilink interface and when it has been traversed through MPLS core random packets may get classified incorrectly by the fabric which leads to packet loss. [PR1223566](#)
- This is a race condition between database creation and database access. Rarely reproducible. There is no functional impact of the core. [PR1225086](#)
- The scale-subscriber license count might increase to an invalid license state with L2TP/LTS clients. This is due to the l2tpd daemon not going through proper state transition on L2TP/LTS clients logout hence license count was not getting updated. The fix will ensure license count is updated on logout regardless of daemon going through proper state transition or not. [PR1233298](#)
- Replacing MPC6E with and ADC based cards causes failure in internal link training. As a result of this ADC based line card will not be booted up normally. [PR1235861](#)
- FPC and RE might stuck in high CPU when DDoS SCFD is turned on. [PR1237486](#)
- Due to a regression issue, presence of errors or traps during ISSU might result in LU/XL based FPC crash. [PR1239304](#)

## Routing Protocols

- In dual REs scenario with NSR and PIM configuration, when backup RE handling mirror updates about PIM received from the master RE, it will delete the PIM session info from its database. But due to a software defect, a leak of 2 memory blocks (8 or 16 byte leaks) will occur for every PIM leave. If the memory is exhausted, the rpd process might crash on backup RE. There is no impact seen on the master RE when the rpd cores on backup. [PR1155778](#)
- Symptoms: With NSR enabled, rpd may core on standby RE when operations like RD modify or RD delete/RD operations are done.(Not always observed). Impact: There is no impact on traffic or other functionality. The core occurs only once on the standby RE. Standby RE recovers completely, with all replication done fully post core. [PR1162665](#)

- When a user executes the "show route advertising-protocol bgp" command for an unconfigured peer, the user may notice that RPD takes much longer to return results than when they run the same command for a configured peer. [PR1173699](#)
- On M120, the bfd of AE interface sticks in "Init" state when using 1:N FEB redundancy. [PR1191217](#)
- I have validated that this is now resolved in 16.1R2 Here are the results when L1 is disabled for Lo0 {master}[edit] labroot@Apollo# run show isis interface IS-IS interface database: Interface L CirID Level 1 DR Level 2 DR L1/L2 Metric lo0.0 3 0x1 Disabled Passive 0/0 Here are the results when L2 is disabled for Lo0 {master} labroot@Apollo> show isis interface IS-IS interface database: Interface L CirID Level 1 DR Level 2 DR L1/L2 Metric lo0.0 3 0x1 Passive Disabled 0/0 [PR1202216](#)
- With nonstop-routing (NSR) enabled, all running protocols include PIM and NG-MVPN will be replicated, if NSR is disabled only under PIM "set protocol pim nonstop-routing disabled", this will remove both PIM and NG-MVPN from replicated list, then adding PIM NSR again by "delete protocol pim nonstop-routing disabled" will not work as expected and PIM will not be added. [PR1203943](#)
- If a NSR enabled router is providing graceful restart support for a restarting peer, and the standby is unconfigured, then rpd may core on the standby during the shutdown. [PR1212683](#)
- RPD leaks memory with the topology and configuration attached to this PR. However it's been confirmed that adding/deleting static flowspec routes in isolation doesn't cause any memory leak. Exact Configuration which causes the leak is still unknown. [PR1213959](#)
- EBGp peer may remain "Idle" at NSR backup-RE, after Interface-down event [PR1215855](#)
- When changing the RD for an existing VRF with established chatty MSDP sessions or deletion/deactivation of MSDP session in the config, the rpd process might crash, which leads to traffic disruption. [PR1216078](#)
- When first multicast packet gets fragments because of bigger in size, the receiver in mvpn scenario does not receive all fragments. The fix of this PR will make sure to wait till last fragment of pim register packet is received at RP before processing the pim resolve request. After last fragment of register packet is received, pim register state is created and pim resolve request is triggered to install multicast route. So, all fragments of the register packet will get forwarded to receiver. [PR1229398](#)
- Juniper implemented BGP4-MIB(including bgpPeerTable and bgpPeerState) per RFC4273. When there's ipv6 bgp neighbor, junos is not able to return correct value for bgp peer. This is caused by bgpPeerTable/bgpPeerEntry is indexed by bgpPeerRemoteAddr, which is SYNTAX'd as IpAddress, a 32bit Interger. But ipv6 address is 128 bit This will cause junos to return 0.0.0.0 which is considered as invalid peer As a workaround, Juniper has defined as extension in "BGP4 V2 MIB" [https://www.juniper.net/documentation/en\\_US/junos15.1/topics/reference/mibs/mib-jnx-bgpmib2.txt](https://www.juniper.net/documentation/en_US/junos15.1/topics/reference/mibs/mib-jnx-bgpmib2.txt) [PR1233790](#)
- When MX router is running protocol BGP, and policy configuration is modified an assertion condition might be hit where the RPD (Routing Protocol Daemon) generates a coredump. [PR1239990](#)

### Services Applications

---

- On MX Series routers with subscriber management feature enabled used as a LAC (L2TP Access Concentrator), a small amount of memory leak is leaked by jl2tpd process on the backup RE when subscriber sessions are logged out. [PR1208111](#)
- In case of massive flapping of subscribers on M120 platform, a memory leak on IQ2E PIC can happen and it can cause inability to attach a configured CoS policer to the newly connected l2tp subscriber [PR1210976](#)
- L2TP subscribers on LNS might get stuck in Terminated state. [PR1215941](#)

### Software Installation and Upgrade

---

- Due to increasing in software requirement and hardware limitation of older hardware; USB installation image may not work correctly in platforms with RE-A-2000 or their variants. The result of using USB install image with these routing engine is for the routing engine to be in a boot loop. [PR1196232](#)

### Subscriber Access Management

---

- Incorrect service-accounting name in radius accounting record if service activated by SRC [PR1206868](#)
- On MX Series routers with subscriber management feature enabled, after GRES switchover "show network-access aaa statistics radius" CLI command display only zeros and "clear network-access aaa statistics radius" doesn't clear statistics as it should. It's a cosmetic issue and communication with Radius server is working fine, the only impact is that affected CLI commands do not work as expected. [PR1208735](#)
- Commit error: "Radius-Flow-Tap LSRI "" is in use by subscriber, cannot be removed from the configuration" might be seen after two consecutive GRES switchovers if a subscriber with lawful intercept mirroring enabled was logged in before the switchovers. [PR1210943](#)
- On MX router with dual REs, after router GRES, if user adds traceoptions filter during GRES not ready period, the authd process might crash. [PR1234395](#)

## User Interface and Configuration

---

- Some configuration objects are not properly handled by dexp. The objects affected can be classified as ?leaf-list? and ?container with presence? in YANG Data Modeling Language (rfc 6020) terminology. ?Leaf-list? can include a list of values (e.g., a list of group names or a list of policy-statement names) but cannot have any other configuration statement nested below it, e.g: - protocols bgp group X apply-groups - protocols bgp group X import - routing-options forwarding-table export ?Container with presence? doesn't convey any additional information and serves just as a present/not\_present flag, e.g: - chassis fpc 7 pic 0 tunnel-services The annotations for statements as above were not processed correctly by dexp (delta-export), were not written to juniper.conf and were shown in ?show | compare? output right after entering configuration mode. [PR1245187](#)

## VPNs

---

- In MVPN mode SPT-only, the first multicast packet is lost when the source is directly connected to the PE. [PR1204425](#)
- On Junos platforms, only VPLS supports automatic-site-id. Configuring automatic-site-id under the L2VPN instance could cause a rpd core. The fix has now been provided to add a commit check to disallow configuring automatic-site-id under a L2VPN instance. With this fix, commit error will be thrown if the user tries to configure automatic-site-id under a L2VPN instance. [PR1214328](#)

- See Also**
- *New and Changed Features*
  - *Changes in Behavior and Syntax*
  - *Known Behavior*
  - *Resolved Issues*
  - *Documentation Updates*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Product Compatibility*

## Resolved Issues

This section lists the issues fixed in Junos OS Release 15.1F7 for the MX Series.

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

- [Resolved Issues: 15.1F7 on page 71](#)
- [Resolved Issues: 15.1F6 on page 92](#)
- [Resolved Issues: 15.1F5 on page 127](#)
- [Resolved Issues: 15.1F4 on page 141](#)

- [Resolved Issues: 15.1F3 on page 147](#)
- [Resolved Issues: 15.1F2 on page 156](#)

### Resolved Issues: 15.1F7

---

- [Class of Service \(CoS\) on page 71](#)
- [Forwarding and Sampling on page 71](#)
- [General Routing on page 72](#)
- [High Availability \(HA\) and Resiliency on page 81](#)
- [Infrastructure on page 82](#)
- [Interfaces and Chassis on page 82](#)
- [Layer 2 Ethernet Services on page 83](#)
- [Multicast on page 84](#)
- [Multiprotocol Label Switching \(MPLS\) on page 84](#)
- [Network Management and Monitoring on page 85](#)
- [Platform and Infrastructure on page 86](#)
- [VPNs on page 91](#)

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

- "show interfaces queue " <if-name> command has three display options: 1. show interfaces queue <if-name> Displays queued/transmitted/dropped packets/bytes for all IFD children. 2. show interfaces queue <if-name> aggregate Displays queued/transmitted/dropped packets/bytes for all IFD children except for IFD RTP traffic 3. show interfaces queue <if-name> remaining Displays queued/transmitted/dropped packets/bytes for IFD RTP traffic only. Note that unlike queued/transmitted/dropped counters, queues depth values cannot be aggregated. With changes done in this PR, the following is true for queues depth values: 1. show interfaces queue <if-name> Displays queues depth values for RTP queues 2. show interfaces queue <if-name> aggregate Displays queues depth values for RTP queues 3. show interfaces queue <if-name> remaining Displays queues depth values for RTP queues. The above logic is the same for physical interfaces, interface-sets and for logical interfaces units.[PR1226558](#)

#### ***Forwarding and Sampling***

- If bandwidth-percent based policer is applied on aggregated Ethernet (AE) bundle without the **shared-bandwidth-policer** configuration statement, traffic will hit the policer even if the traffic is not exceeding the configured bandwidth. [PR1125071](#)
- Firewall module (daemon dfwd) on Routing Engine always leaks some memory upon config commit with the following configurations: **set routing-options forwarding-table export qos3 set policy-options policy-statement qos3 term 1 from source-address-filter 148.10.200.3/32 exact set policy-options policy-statement qos3 term 1 then forwarding-class Q0-class**. The issue can be seen in both high availability (HA) and standalone instances. [PR1157714](#)

- Commit gives the following error when apply-groups is configured under the bridge domain: error: Check-out failed for Firewall process (/usr/sbin/dfwd) without details. [PR1166537](#)
- The Sampling Route-Record Daemon (SRRD) process does not delete routes when the DELETE is received from the routing protocol process (rpd) in a few configuration cases. This results in buildup of memory in the SRRD daemon. Once SRRD reaches the limit, it crashes and restarts itself. This happens only when one certain family is not configured on all of the FPC clients (such as an FPC with inline J-Flow enabled or a PIC with PIC-based sampling enabled in one client). For example, only the IPv4 family is configured in all the clients, and IPv6 and MPLS families are not configured for sampling in any of the clients. [PR1180158](#)
- The changes to the Sampling Route-Record Daemon (SRRD) sampling route reflector daemon - new architecture for sampling) process between Junos OS Release 14.2R5.8 and 14.2R6.5 severely reduce MX80 Series available memory and therefore Routing Information Base (RIB) and Forwarding Information Base (FIB) scaling. [PR1187721](#)
- Starting with Junos OS Release 14.2R1, FPC offline could trigger Sampling Route Record (SRRD) daemon restart. [PR1191010](#)
- If a two-color policer is configured on MX Series routers with certain MPCs, more traffic than the limited traffic might be passed when packet size is less than 128 bytes. [PR1207810](#)
- On MX Series platforms with Enhanced Subscriber Management mode, if default forwarding classes are referenced by subscriber filters, commit configuration changes after Graceful Routing Engine Switchover (GRES) will fail. [PR1214040](#)
- Use of the firewall filter family "any" with shared-bandwidth-policer on an MC-AE interface does not reconfigure bandwidth or divide the policer. When standby becomes active after A/S switchover, all packets are dropped. [PR1232607](#)
- On MX Series devices with "ipv4-flow-table-size" or "ipv6-flow-table-size" config, if the sampling instance is not defined under the chassis hierarchy (sampling instance is not associated to FPC), after rebooting the device, the "ipv4-flow-table-size" or "ipv6-flow-table-size" does not propagate to the FPC. [PR1234905](#)

### **General Routing**

- The %Count column is replaced with %Memory to give a better view of resource utilization. [PR1040789](#)
- The wrong byte count was seen in the ipfix exported statistics packets for MPLS flows. [PR1067084](#)
- Certain VTY JNH commands on Trinity platform will not decode properly. [PR1094955](#)
- During initial ramp-up of an IPsec session, a race condition might cause the mspmand process to crash in rare circumstances. [PR1116487](#)
- On MX Series platforms with MS-MPC/MS-MIC in use, if the NAT session is freed or removed without removing timer wheel entry, MS-MPC/MS-MIC might crash. It is a timing issue where just before invoking the timer wheel callback, the NAT session extension got freed or removed. [PR1117662](#)



- On MX Series platforms, the MS-MPC might crash. The exact trigger of the issue is unknown; normally, this issue might happen over long hours (for example, within a week) of traffic run (for example, running HTTP/HTTPS/DNS/RTSP/TFTP/FTP traffic profile). Core files might occur with this message: the composite next hop.**Program terminated with signal 4, Illegal instruction**[PR1124466](#)
- The rpd might crash when local-switching is configured with connection-protection for L2Circuit. This problem only occurs after setting composite next hop for L2Circuit using **set routing-options forwarding-table chained-composite-next-hop ingress l2ckt**. [PR1129940](#)
- The jsscd might crash in a static-subscribers scaling environment (for example, 1,12,000 total subscribers, 77,000 DHCP subscribers, 3,000 static-subscribers, 32,000 dynamic VLANs), and the subscribers might be lost. [PR1133780](#)
- On MX Series platforms, the Max Power Consumption of MPC Type 13D (model number: MX-MPC1-3D) would exceed the default value because of a software issue. For example, the value might be shown as 368 watts instead of 239 watts when max ambient temperature is 55 degree Celsius. [PR1137925](#)
- If a line card crashes early in a unified ISSU warm boot, the CLI might report ISSU success, resulting in a "silent ISSU failure". [PR1154638](#)
- The speed command **auto-10m-100m** enables you to autonegotiate the speed maximum to 100 mbps. [PR1155196](#)
- The rpd might crash after EVPN was configured when extra bits in the ESI label extended community were set in addition to the single-active bit [PR1158195](#)
- Software OS thread on the line card is doing a busy loop by reading the clock directly from hardware. Sometimes it seems the thread is getting incorrect values from the hardware register and thus waits in the busy loop. After the busy loop occurs for a certain time period, the line card crashes and reboots. This is a rare condition. [PR1160452](#)
- Upon receiving some specific packets, a compute CPU of the MS-MIC/MS-MPC will stop refreshing the inactivity-timeout of sessions despite receiving traffic matching them. As a result, an affected session will be removed when the inactivity-timeout time has expired. [PR1161040](#)
- The default per-packet load balancing PPLB export policy created for Ethernet VPN (EVPN) has been removed from Junos OS. It was used to enable per packet load balancing for EVPN routes on certain MX Series platforms but not all. Now per-packet load balancing needs to be configured explicitly. [PR1162433](#)
- Interfaces routing status message **xxx.xxx.xxx.xxx <Up Broadcast Multicast Localup>** might be reported on an interface that is not associated with the configuration change, such as a bridge-domain addition. It should be reported only if there is any change in the IFL parameters. This is an info(6) level message for debug purpose, so we can safely ignore the cosmetic problem. **rpdd[xxx]: %DAEMON-6: EVENT Flags ge-1/0/4.0 index 371 10.180.230.8/24 -> 10.180.230.255 <Up Broadcast Multicast Localup> rpdd[xxx]: %DAEMON-6: EVENT Flags irb.110 index 326 10.9.17.254/22 -> 10.9.17.255 <Up Broadcast Multicast Localup> rpdd[xxx]: %DAEMON-6: EVENT Flags irb.190 index 373 10.9.53.254/22 -> 10.9.53.255 <Up Broadcast Multicast Localup> .** [PR1162699](#)

- On MX Series routers with services PIC (MS-DPC/MS-MPC/MS-MIC), the ICMP time exceeded error packet is not generated on an IPsec router on the de-encapsulated side. [PR1163472](#)
- When the MS-MIC or MS-MPC installed in an MX Series router is processing traffic, and the IPsec policy configuration is changed by means of adding or updating a policy, mspmand process crash might occur. [PR1166642](#)
- The sampled process continues logging events in the traceoption file after the traceoption for the sampled process is deactivated. This issue can occur if there is no configuration under **forwarding-options sampling** but another configuration for sampled is present (for example, port-mirroring). [PR1168666](#)
- When you enable LLDP and the interface description is long (more than 32 characters) on the remote switch, the layer2 Control Protocol Process (l2cpd) might crash with a core file if you are performing an SNMP MIB walk because LLDP code is running within l2cpd. [PR1169252](#)
- When MS-MPC is used, if any bridging-domain-related configuration exists (such as family bridge, vlan-bridge, or family evpn) in some cases, continuous MS-MPC crash and traffic loss might occur. [PR1169508](#)
- When upgrading or rebooting the router, the following logs might be seen in Junos OS Release 15.1F5. There is no impact and they can be ignored. The logs are seen because agentd is trying to read the forwarding class entries at system boot time too early, when they are not yet created. [PR1173137](#)
- When maximum-ecmp 64 is enabled and if an IS-IS route has many next hops or above the maximum, rpd might crash because the next hop gateway addresses are getting overwritten and stored in a circular buffer. Note: In the worst case (if all the next hops are IPv6), only 38 ECMP next hops are fully supported for IS-IS IPv6 instead of 64. [PR1174892](#)
- Changing "inline-services flow-table-size" might cause memory-related errors to be logged until the FPC is rebooted. [PR1176186](#)
- On MX Series platforms with MPC7E\MPC8E\MPC9E cards, if line rate traffic in which packet size is 512 bytes on 100G interfaces for full mesh mode, some traffic might be dropped. [PR1176706](#)
- On dual Routing Engine system, if the master Routing Engine is running Junos OS Release 13.3R9/14.1R7/14.2R5/15.1R3/16.1R1 or later, and the backup Routing Engine is running a Junos OS release prior to 13.3R9/14.1R7/14.2R5/15.1R3/16.1R1, a major alarm is raised. This is cosmetic and can be safely ignored. [PR1177571](#)
- CGNAT-NAT64: A few port leaks are observed for the EIM/EIF IPv4 traffic (2M sessions) from the public side. [PR1177679](#)
- When the Switch Interface Board (SIB) is pulled excessively, the interface might become wedged, causing permanent drop. [PR1177753](#)
- On dual Routing Engine platforms, if interface changes occur on the Aggregate Ethernet (AE) interface that result in marking ARP routes as down on the AE interface (for example, bringing down one of the member links), because of an interface state pending operation issue on the backup RE, the backup Routing Engine might crash and reboot

with an error message such as the following: **panic:rn timer\_index\_alloc: nhindex XXX could not be allocated err=X**. [PR1179732](#)

- If the MIC-3D-4XGE-XFP is used with MPC2E-3D-NG or MPC3E-3D-NG, the interfaces on the MIC-3D-4XGE-XFP connected to a DWDM device might flap continuously. [PR1180890](#)
- In the CGNAT CLI **show service alg conversations** fails to display parent session status for ALG conversations. [PR1181140](#)
- The Packet Forwarding Engine ukern MDI thread is hogging CPU. MDI feature is not supported on the MX86 platform. [PR1181172](#)
- In case of point-to-point interfaces and unnumbered interfaces rpd crash might be seen in corner cases on configuration changes. [PR1181332](#)
- Ping fail in lt interface ip after Load baseline config & Rollback to mvng [PR1181517](#)
- After doing GRES, some fabric planes might go to check state on MX2020/MX2010 platform with MPC7E/MPC8E/MPC9E. [PR1182851](#)
- With NAT translation-type as napt-44, a few sessions are getting stuck upon deactivating/activating service-set or corresponding applications with traffic running. The same symptom is seen upon deactivating/activating service-set with traffic running and with 'deterministic-napt44' translation type as well. [PR1183193](#)
- CGNAT Pool stats for "Available address" is shown incorrect for destination pool. The available address shows zero even though destination NAT IPs are available [PR1183538](#)
- With BGP add-path and consistent-hash enabled, when a BGP-learned route prefix with multiple paths (next hop) is installed in the forwarding-table, all the next hops should be reachable/resolvable at the time of installing the route in the forwarding-table. However, if any of the next hops are not resolvable at that time, incorrect route programming will occur on the Packet Forwarding Engine. In this case, traffic forwarded to this prefix will be affected. [PR1184504](#)
- When an IPv4 firewall filter has 2625/32 destination in the prefixlist, the filter attached to subscriber interface gets broken. [PR1184543](#)
- Continuous reporting of the following messages might be noticed sometimes while bringing up all IFD/IFL/IFF states at once. In scale scenario the size of this information might exceed the available demon rlimit memory availability. [PR1184948](#)
- In an IPv6 environment, when you add a link local neighbor entry on the subscriber interface then add a new lo0 address, if you delete this neighbor entry and the subscriber interface, the next hop information is not cleaned properly. As a result, the rpd process might crash. The routing protocols are impacted, and traffic disruption will be seen due to loss of routing information. [PR1185482](#)
- On MX Series platforms with MS-MICs or MS-MPCs installed, when the ams-interface is configured in warm-standby mode without adding any members, configuration commit will lead to rdd crash. [PR1185702](#)
- In IPv6 environment with graceful Routing Engine switchover (GRES) enabled, when a new prefix (global address) is added on the donor interface (in this case, loopback

interface), and then GRES is performed, the ksyncd process might crash because of a kernel replication error. [PR1186317](#)

- Traffic destined to Virtual Router Redundancy Protocol (VRRP) Virtual IP (VIP) address or transit traffic with the destination MAC address as VRRP Virtual Media access Control (VMAC) that has payload beyond 166 bytes (excluding headers) is dropped because "my-mac check failed" on MPC7E/8E/9E [PR1186537](#)
- After loading CoS related configuration on an MPC5E/MPC6E/MPC2E-NG/MPC3E-NG linecard, these error messages might be seen: "trinity\_insert\_ifl\_channel:6449 ifl 495 chan\_index 495 NOENT" "jnh\_ifl\_topo\_handler\_pfe(11591): ifl=495 err=1 updating channel table nexthop" [PR1186645](#)
- CHASSISD\_I2CS\_READBACK\_ERROR errors might occur on single occurrence of I2C read failure. These are transient errors. The errors might be seen randomly without any particular trigger. After the fix, these messages are seen only when there are three consecutive I2C read failures. [PR1187421](#)
- By default SNMP will cache SNMP values for 5 seconds. Sometimes the kernel will cache these values for a longer duration. [PR1188116](#)
- On MX series with MS-MIC installed, if dynamic routing protocols are configured over IPsec, sometimes the MS-MIC might crash. [PR1188275](#)
- On MX Series platforms with Junos OS Release 15.1 or later, the, LLDP PDU gets dropped on the internal Ethernet (FXP) interface. [PR1188342](#)
- On MX routers, a vulnerability in IPv6 processing might allow a specially crafted IPv6 Neighbor Discovery (ND) packet to be accepted by the router rather than discarded. The crafted packet, destined to the router, will then be processed by the Routing Engine (RE). A malicious network-based packet flood, sourced from beyond the local broadcast domain, can cause the Routing Engine CPU to spike, or cause the DDoS protection ARP protocol group policer to engage. When this happens, the DDoS policer might start dropping legitimate IPv6 neighbors as legitimate ND times out. Refer to [JSA10749](#) for more information. [PR1188939](#)
- Ingress queuing configuration on next-generation MPC2E is leading to host loopback wedge. [PR1189800](#)
- On MX240/MX480/MX960/MX2010/MX2020 platform, in rare cases, the MPC4 line card might never come back online after you reboot the chassis using the **request system reboot both-routing-engine** command. [PR1190418](#)
- SSH keys are not preserved across upgrades. [PR1190852](#)
- On MX2010/MX2020 platform with Enhanced Switch Fabric Boards (SFB2) installed, when Link Fault Management (LFM) session is configured on MPC7/8/9 and with timeout of 300 ms or less, it might flap during another MPC's offline sequence. [PR1191546](#)
- When polling an si-interface hosted on a next-generation MPC Non-HQoS line card (MPC2E-3D-NG, MPC3E-3D-NG) ,a 10-second delay occurs that might break SNMP polling. [PR1192080](#)

- As described in RFC 7130, when LACP is used and considers the member link to be ready to forward traffic, the member link MUST NOT be used by the load balancer until all the micro-BFD sessions of the particular member link are in Up state. [PR1192161](#)
- If a message received from an LLDP neighbor contains "Port Id" TLV that has "Interface alias" subtype and is longer than 34 bytes, subsequent running of "show lldp neighbors" might lead to l2cpd crash. [PR1192871](#)
- On MX series with MPC3/MPC4/MPC5/MPC6 (might also affect EX92xx, for example, with EX9200-32XS, SRX5000 line platforms), the VSC8248 firmware on the MPC crashes occasionally. This PR enhances the existing VSC8248 PHY firmware crash detection and recovery, helping recover from a few corner cases where the existing Junos OS workaround does not work. [PR1192914](#)
- From Junos OS release 15.1F5/16.1R1 and above, when the forwarding-table export policy with install-next hop is configured, route resolution might be incorrect when the forwarding next hop of the dependent route changes. It might cause incorrect LSP or even an entropy label to be installed for forwarding, which results in packet loss. [PR1193731](#)
- MAC routes received from the control plane are not installed in the EVPN mac-table [PR1193754](#)
- With GRES (graceful-switchover) and nonstop-bridging configured in Juniper devices with dual Routing Engines, the backup Routing Engine might run into high CPU usage due to abnormally high CPU utilization by firewall daemon. The abnormally high CPU usage might impact the functions that backup Routing Engine works for. [PR1193891](#)
- In port mirror, IPv4 inbound traffic might not get mirrored to the 10G Analyzer interfaces types. [PR1194139](#)
- When Multicast-only Fast Reroute (MoFRR) activated, multicast source route flapping leads to corresponding multicast traffic 100 percent drop. [PR1194730](#)
- Precision Time Protocol (PTP) support for next-generation MPC2E and MPC3E-NG is not working in Junos OS Release 16.1R1. This issues occurred because incorrect branch synchronization removed the support checks for these MPCs. [PR1194734](#)
- On platforms running Junos OS with FreeBSD10, if tracing is enabled, because the log file pointer is not being handled correctly for log file rotation, the rpd process might crash when the log file rotates. [PR1196318](#)
- Distributed BFD session using inline-redirection on MX Series Virtual Chassis might not work if the ANCHOR PFE is not within the same chassis member as the interface where the BFD packet is received from the peer devices. [PR1197634](#)
- L2VPNs or L2Circuit services along with lengthy interfaces descriptions might lead to memory leak in variable-sized malloc block, which in turn results in rpd crash due to "out of memory". [PR1198165](#)
- Continuous error messages are generated during 2X100GE CFP2 OTN MIC online on MX2K. This error message means PCI control signal communication failure between Packet Forwarding Engine on MPC6E and PMC Sierra OTN framer (pm544x) on MIC 2X100GE CFP2 OTN. [PR1198295](#)

- With MPC8/9 MRATE MIC with plug-in optics module (QSFP28-100GBASE-LR4), bit errors might be seen. [PR1200010](#)
- With MPC-NG or MPC5E hardware, the range of the queue weights on an interface is from 0 to 124. Because every queue must have an integer value of queue weight, it might be impossible to assign the weights in exact proportions to the configured transmit-rate percentage. Therefore, when a physical interface operates in a PIR-only mode, this might cause imprecise scheduling results. [PR1200013](#)
- On MX Series platforms, the mspmand process might crash on the MS-MPC with XLP B2 chip (for example, REV17). The exact trigger is unknown. It is usually seen with 70% to 90+% CPU load conditions. [PR1200149](#)
- GUMEM errors for the same address might continually be logged if a parity error occurs in a locked location in GUMEM. These messages should not impact performance. The parity error in the locked location can be cleared by rebooting the FPC. [PR1200503](#)
- The MSPMAND might crash when an encrypted packet is received out of the range of replay-window size. The issue might occur in peak loads whereby encrypted packets received, out of order cause drops in the network. [PR1200739](#)
- SFB2: MPC9 observed ~80G packet drop (20G x 4 PFEs) with offline/online 8th plane with 1.6Tbps line rate traffic (400G x 4 PFEs). Full line rate traffic resumed after fabric training was completed. This issue is not seen if traffic load is less than 1520G (380G per PFE). [PR1201238](#)
- A dynamic tunnel gets timed out every 15 mins by default, and then re-tries to create another tunnel. This happens if the route obtained from IGP is non-forwarding. [PR1202926](#)
- The Packet Forwarding Engine might install a next hop incorrectly and cause traffic loss, if there is a next-hop policy pointing to an IPv6 address that needs to be resolved. [PR1204653](#)
- On MX240/MX480/MX960 platform with RE-S-2000, the hard drive information on the Routine Engine is missing in "show chassis hardware detail" output after upgrading to Junos OS Release 15.1 and later. This is just a display issue and this has no impact on any functionality. [PR1205004](#)
- Problem - In case of local source and with any-source multicast (ASM) multicast-only Fast Reroute (MoFRR) enabled, the default Multicast Distribution Tree (MDT) traffic loops back to the originating router on the MoFRR backup interface, thereby causing continuous IIF\_mismatches. With the current MoFRR code, because the source is local, SPT BIT is set by default. Therefore an (S,G,rpt) PRUNE is sent out of the MoFRR active interface. But an (S,G,rpt) PRUNE is not sent out of the MoFRR backup interface (missing code). [PR1206121](#)
- When Path Computation Element Protocol (PCEP) is enabled and label-switched paths (LSPs) are undergoing changes, for example, make before break (MBB) for rerouting, the rpd sends those updates to the PCE. However, when the PCEP session to PCE goes down, these updates are canceled, but the rpd fails to completely reclaim the memory allocated for these updates. This requires more rpd memory every time the connection to PCE goes down while LSPs are simultaneously going through MBB changes. This issue will be especially noticeable when connectivity to PCE goes UP

and DOWN continuously. If the connection is in steady state, either UP or DOWN, then the memory leak will not happen. [PR1206324](#)

- When FPC software reads watchdog timer on boot-cpld register, sometimes it gets unexpected value "0x0000" because it was not refreshed by stroker in time. As a result of this "npc\_check\_boot\_cpld: boot cpld watchdog time access error @0xff000005, expected 0x0faf (4015) got 0x0000" errors could be seen in the system log. [PR1206624](#)
- The l2ald might thrash when the targeted-broadcast is configured on the EVPN Integrated Routing and Bridging (IRB) interface. [PR1206979](#)
- When an egress Packet Forwarding Engine (NG-MPC3E) is oversubscribed, it applies flow control to the ingress Packet Forwarding Engine (MPC7E). - The fabric delay buffer memory utilization on the ingress PFE (MPC7E) goes up because the flow control from the egress PFE. The default WRED drop profile for the low-priority fabric queues does not aggressively drop the low-priority traffic. Have separate default WRED drop profiles for low and high-priority fabric queues. Set up the default WRED drop profile for the low-priority queues to drop the traffic more aggressively so that high-priority traffic can be protected. [PR1207417](#)
- This is a rare race condition of multiple interrupts not being handled properly on MX Series platforms with MPC7E/MPC8E/MPC9E and PTX platforms with FPC3-PTX-U2/FPC3-PTX-U3 that can lead to core files.. It is hard to reproduce. The interrupt code is optimized to avoid the unnecessary call to prevent the issue. [PR1208536](#)
- When using the **show chassis hardware detail** command in Junos OS Release 15.1 or later to display chassis components, the flash memory disk and hard disk serial numbers might be truncated to 15 characters. [PR1209181](#)
- The logic to calculate the IPsec phase2 soft lifetime has been changed in Junos OS Release 14.2R6, resulting in an interop issue in certain scenarios. A hidden command is provided as part of this PR that will revert the soft lifetime logic to that used in Junos OS Release 11.4 release. [PR1209883](#)
- BGP PIC installs multiple MPLS LSP next hops as active instead of standby in PFE. This can cause a routing loop. [PR1209907](#)
- On MX Series platform, if any inline feature is configured (such as inline BFD, CFM ,or PPP), the FPC might crash and core files are generated. [PR1210060](#)
- Puppet/Chef functionality is not supported on QFX5110 in the current release. [PR1210477](#)
- When inline J-flow is enabled, the flow sequence number in the flow data template is set to zero on MPC5E/6E/7E/8E/9E and MPC2E-NG/MPC3E-NG while exporting the flow record to the collector. Depending on the implementation, certain collectors might result in collector fail to decode the flow record and missing flows. [PR1211520](#)
- FPC major alarm and "MQSS overflow" error messages might be reported on MPC9E running at line rate with small packet sizes. This issue causes no traffic loss. [PR1213391](#)
- On detection of a zero-length memory allocation in the SDB database, a forced rpd crash would be seen. [PR1215438](#)
- This issue happens only with RLT configuration and only on 16.1 and beyond. [PR1216991](#)



- Kernel crash and router reboot might happen when committing RLT configuration. [PR1218326](#)
- If RS/RA messages were received through an ICL-enabled(MC-AE) IFL, packet loss would be seen and last for a while. [PR1219569](#)
- On MX Series, when you enable the feature VRRP delegate-processing ae-irb, VRRP and BFD might be flapping. [PR1219882](#)
- The master CB/RE offline or OIR could lead the pcie link down between SFB2 and CB during link reset. As a result, some of SFB could be in check status followed by fabric healing. With the fix, the software will try 5 times to help the graceful link come up. When the issue occurred, following chassis alarm could be seen: "Minor Check plane <idx> Fabric Chip" where <idx> is SFB slot number. [PR1219890](#)
- On turning the circuit breaker on or off for a DC PEM multiple times, on certain occasions, the craft interface shows PS LED as RED even when that DC PEM is online and green with no alarms. There is no operational impact, and it is not seen every time when the activity is repeated. [PR1220693](#)
- When fpc-pfe-liveness-check is configured, Packet Forwarding Engine liveness detection might incorrectly report a PFE failure event under a severe interface congestion situation. [PR1220740](#)
- On MX Series Virtual Chassis partial or complete traffic loss for streams via AE interfaces might be observed in certain scenarios. For example, if vcp ports were de-configured and re-configured again, the two consecutive global GRES switchovers were performed, and the MPC hosting AE child links was reloaded, traffic loss would be observed after the MPC booted up. This issue occurs because of incorrect programming of the AE interface on its Packet Forwarding Engine. [PR1220934](#)
- On MX series with "pppoe dynamic-profile and service-name-table xx" configured, when you configure the prefix or any interface configuration and after commit the configuration, the output of "show pppoe service-name-tables xx" is displayed as "Service Name Table not found: xx". [PR1221278](#)
- After Junos OS Release 15.1, the behavior of storage devices enumeration in kernel level has been changed. Device enumeration in legacy Junos OS (pre 15.1) will show CF and Disk as ad0 and ad1, respectively. Device enumeration after Junos OS Release 15.1 will show CF and Disk as ad1 and ad0 instead in the result of "show chassis hardware". This might be inconsistent for other result of output, such as "show system boot-messages" and "show log messages". [PR1222330](#)
- Due to a defect related to autonegotiation in a Packet Forwarding Engine driver, making any configuration change to an interface in MIC "3D 20x 1GE(LAN)-E,SFP" might lead to interface flapping. [PR1222658](#)
- When you receive alignment errors on a 10-Gigabit Ethernet port, the MAC control frames counter might show a huge value. [PR1224632](#)
- On the MX2020 router, when all the Switch fabric Boards (SFBs) are removed, there is no available fabric in the system, but FPCs remain in online state. There is no problem in offlining these SFB/SFb2s. [PR1227342](#)



- Flowstat reply has incorrect DL type. In the following output, the flowstat reply shows DL type as 0xcc88 instead of 0x88cc : user@host> show openflow flows detail Flow name: flow-65536 Table ID: 1 Flow ID: 65536 Priority: 32768 Idle timeout(in sec):0 Hard timeout(in sec): 0 Cookie: 1 Match: Input port: wildcard Ethernet src addr: wildcard Ethernet dst addr: wildcard Input vlan id: 50 Input VLAN priority: wildcard Ether type: 0x88cc IP ToS: wildcard IP protocol: wildcard IPv4 src addr: NA IPv4 dst addr: NA IPv6 src addr: NA IPv6 dst addr: NA ICMPv4 type: wildcard ICMPv4 code: wildcard Source port: wildcard Destination port: wildcard Action: Output port CONTROLLER, user@host> [PR1228383](#)
- The Random Load Balancing feature does not function; all traffic goes to one of the load-shared egress links instead of being shared across all the links. [PR1230272](#)
- On XQ-based line cards, in a rare condition, if offline/online the FPC or link flap, some error messages might be seen. [PR1232686](#)
- When a Switch Fabric Board (SFB) is offline/onlined, the HSLink crc error values are not cleared properly, which triggers an unexpected Link error/ SFB check alarm for another SFB. [PR1234224](#)
- After the backup Routing Engine is replaced, the new backup Routing Engine cannot synchronize with master Routing Engine if **dynamic-profile-options versioning** is configured. This is because the code checks if any dynamic profile is configured before enabling dynamic-profile-options versioning. If so, it throws a commit error. But there is no need to check when the Routing Engine is in backup state. [PR1234453](#)
- When PIC-based MPLS J-flow is configured and MPLS packets are being sampled at egress (to be sent to service pic), the sampled packets do not reach a service pic. As a result, no MPLS J-flow flows get created. [PR1236892](#)
- When an MPC6E slot#10 is installed in an SFB2-based MX2020 router, and an SFB#4~7 is offlined/onlined once, the next slot SFB will get an 'SFB check alarm' unexpectedly. For instance, SFB#4 offline/online triggers an SFB#5 check alarm. [PR1237134](#)
- Because of the lack of proper boundary checks in code, the MS-MPC might crash when receiving internally corrupted frames from other FPCs that have hardware failures or incorrect rewrite programs. [PR1237667](#)
- In a BGP-PIC scenario involving a change in the IGP topology, (for example, a link failure in the IGP path), causes traffic outage for certain prefixes. This issue occurs because the unilist next hops for these prefixes are in a broken state. [PR1239357](#)
- Routes learned over external BGP (EBGP) multipath peering might not get installed in the forwarding table. As a result, traffic black-holing for the affected destinations is silently dropped or discarded. This will only happen if in addition to EBGP multipath there is also a multi-hop command enabled for that peering and a unicast Reverse Path Forwarding (unicast RPF) check is enabled over the involved interfaces. Corresponding routes will end up stuck in the KRT queue and related KRT log messages containing error code **EINVAL -- Bad parameter in request** are seen in the logs. [PR1241501](#)

### **High Availability (HA) and Resiliency**

- Unified ISSU will fail when trying to upgrade from 15.1F6 to 16.1 images and higher. [PR1187779](#)

- On all platforms, during unified ISSU, the connection between the master Routing Engine and the backup Routing Engine might be broken. [PR1234196](#)

### ***Infrastructure***

- From Junos OS Release 15.1 onward, smartd error message of Unigen SSD might be seen. The smartd process reads SSD attributes and checks on 197-current-uncorrectable, 198-offline-uncorrectable; by default. To Unigen, 198 is not = Offline-Uncorrectable, it is 'Total Count of Read Sectors'. Because it is Total-Read, attribute(198) always carries values, and smartd reports it as 'Offline Uncorrectable Error'. [PR1187389](#)
- The statistics information of em0 is 0 when checked by SNMP or the CLI **show** command. [PR1188103](#)
- Polling SNMP QoS queue statistics along with physical interface statistics might result in flat values for QoS queue statistics. The flat values could give a false impression that spikes are happening in the queues. [PR1226781](#)

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

- Clarify that BER is reported over a 5-second interval. This differs from Q-factor, which is reported over a 1- second interval. [PR1159027](#)
- In a subscriber management environment, the "jpppd: RLIMIT\_STACK & RLIMIT\_SBSIZE messages" are marked incorrectly at NOTICE level instead of at INFO level. The message is not indicating a memory leak; it is only a system log message that should be suppressed, and it is only used for debugging purpose. This issue might be seen even if PPP/PPPoE is not enabled. [PR1178895](#)
- In the hsl2 toolkit, there is a process that periodically checks the ASICs that communicate through it. Due to a bug in the toolkit code, the process used invalidates the very ASIC that it used to process, and a crash occurs. [PR1180010](#)
- When there is a configuration change about OAM CFM, cfmd memory leak is observed and sometimes also might trigger cfmd crash and show the following messages::  
/kernel:Process (44128,cfmd) has exceeded 85% of RLIMIT\_DATA: used 378212 KB Max 393216 KB. [PR1186694](#)
- The jpppd might crash with a core file because of a memory heap violation associated with processing Multilink Point-to-Point Protocol (MLPPP) requests [PR1187558](#)
- If a "filter" command is present in a Point-to-Point Protocol over Ethernet (PPPoE) traceoptions configuration, the resulting log file will contain only a few of the messages specific to establishing to the PPPoE session. However, the log file will contain information related to other sessions established at the moment. [PR1187845](#)
- In an Operation, Administration, and Maintenance (OAM) Connectivity Fault Management (CFM) scenario on AE interfaces with maintenance-domain level (for example, 3) configuration, when OAM CFM LBM messages with a level smaller than the configured level are sent to the ingress interface of VPWS with QinQ encapsulation, they are not dropped by the ingress PE. [PR1191818](#)

- On an MX Series Virtual Chassis, CFM sessions on the ae interface are not distributed to the FPC when member-1 chassis are chosen as primary. [PR1198447](#)
- MAC addresses are incorrectly assigned to interfaces by the MX Series Virtual Chassis SCC (global) chassisd daemon, leading to duplicate addresses for adjacent FPCs. [PR1202022](#)
- A CFMD core file will be generated upon commit if \* CFM is configured and the ICC format is incorrectly configured for MCC (for example, the ICC name-format does not start with a character). [PR1202464](#)
- For the duration of GRES, if an asynchronous message for RTTABLE is received at DCD during initialization, it might result in unexpected state changes, and the traffic forwarding might be affected. This is a timing issue, it is difficult to reproduce. [PR1203887](#)
- In very rare conditions, FPC might crash when the CLI command **request chassis mic offline fpc-slot <fpc-slot> mic-slot <mic-slot>** or **request chassis pic offline fpc-slot <fpc-slot> pic-slot <pic-slot>** is executed. This is due to a software defect in which the SFP diagnostics polling function tries to access already destroyed SFP data structure by MIC/PIC offline. [PR1204485](#)
- When configuring "vlan-tags" for any interface, if the interface configuration is changed continually, the dcd process might experience leak. If the memory is exhausted, the dcd process might crash. [PR1207233](#)
- The **show interfaces terse routing-instance all** command has wrong display format when there are multiple addresses. [PR1207272](#)
- On Junos OS 14.2 and later releases, if asymmetric-hold-time, delegate-processing, and preempt hold-time are configured, when the neighbor's interface comes up again, the "asymmetric-hold-time" feature cannot be used as expected. [PR1219757](#)

### **Layer 2 Ethernet Services**

- This issue occurs when running LACP between Juniper and Cisco devices with different timers (Juniper fast and Cisco slow) on both sides. On the Cisco side it takes almost 90 sec to bring the interface down from the bundle. When one interface is removed from the LAG on the Juniper side, the lead to the Cisco side to time out to bring the interface down from bundle. This results in unexpected behavior outage on the network. [PR1169358](#)
- A new static MAC is configured under AE interface, but the MAC of the LACP PDUs sent out is not changed. [PR1204895](#)
- When Graceful Routing Engine Switchover (GRES) is enabled, after Routing Engine switchover, the local MAC address is not learned anymore from the local CE device in the VPLS instance because of spanning-tree "discarding" in the kernel table. [PR1205373](#)
- On MX Series platforms, if chassis- level configuration is used to offline the FPC after detecting major errors, FPC will be offlined. But if the committing configuration is performed after offlining the FPC, the FPC will be brought online back again. [PR1218304](#)
- During the unified ISSU process, if the first ISSU is aborted for some reason, an internal timer will not be cleaned up, and the new lacpd will be forked up. This causes the

second ISSU in the backup Routing Engine to be aborted in the daemon prepare phase. It will not proceed further. [PR1225523](#)

- When DPC cards are used and the **set chassis fpc-pfe-liveness-check** command is configured, some alarms can be seen on the DPC cards ( for example, "/var partition is full) during upgrading to 15.1F2-S13 from 15.1F2-S12. When trying to downgrade to 15.1F2-S12, and the alarm is cleared, and when upgrading to 15.1F2-S13, the alarm is seen again. [PR1237218](#)

### ***Multicast***

- The rpd creates an indirect next hop when a multicast route (S,G) needs to be installed when listeners show their interest to S,G traffic. The kernel would then creates a composite NH. In this case, it appears to be a P2MP MCNH that gets created. When any member interface is not a Packet Forwarding Engine specific interface (such as, Vt, LSI, IRB or any other pseudo interfaces), the kernel throws messages indicating that FMBB cannot be supported. These messages are harmless and does not have any impact. [PR1230465](#)

### ***Multiprotocol Label Switching (MPLS)***

- Up until version -10 of the BGP-LS draft, the OSPF DR node representation was ambiguous. One could represent DR node as 'AdvertisingRouterId-InterfaceIpAddress' or 'InterfaceIpAddress-1'. JUNOS used to follow 'InterfaceIpAddress-1' format. Starting version '-11' of the BGP-LS draft, the representation for OSPF DR node must be 'AdvertisingRouterId-InterfaceIpaddress'. Junos OS now follows the latest format. [PR1085219](#)
- User is allowed to configure both "load-balance-label-capability" and "no-load-balance-label-capability" together. This is incorrect and confusing. [PR1126439](#)
- In the penultimate-hop popping (PHP) link goes down and the other router becomes MP for an LSP, the next-hop link for the same LSP, the next-hop link for the same LSP goes down. The router becomes point of local repair (PLR) for the same LSP. Thus effectively the router is both MP and PLR for the same LSP. In this scenario, the router sends an incorrect PathErr message for the backup PSB. It sends the Bad strict route PathErr instead of the Tunnel local repaired PathErr. [PR1132641](#)
- When PCE updates the delegated LSP, no-install-to-address configured under the LSP stanza is not honored. [PR1169889](#)
- RSVP-signalled p2mp sub-LSP with at least 1 or more sub-LSPs in a down state might not get re-optimized in the event of a transit core link going down. If there are no sub-LSPs in a down state at the time of re-optimization, this issue will not be seen. This can cause traffic drop over the sub-LSPs that are carrying traffic, but are unable to get re-optimized. [PR1174679](#)
- Multiple RLFA backup gateways (one using spring inner label and the other using TLDP label) can get programmed if the given node is the PQnode to another node in the network that does not use SPRING RLFA backup for its LDP route, resulting in ECMP among backup nexthops. Semantically, both the gateways provide the same protection path and the TLDP- based gateway is preventing the sanity check of the SPRING backup path. [PR1176489](#)

- On Juniper devices with "link-protection" configured and with "optimize-adaptive-teardown p2p" configured, rpd might crash after link flap. [PR1186003](#)
- Packets will be out of order if they are generated by the Routing Engine and go over unicast/ECMP. [PR1193697](#)
- Changing the configuration under both [ protocols pcep ] and [ protocols mpls lsp-external-controller ] might trigger rpd to crash because of a race condition. [PR1194068](#)
- In L3vpn with chained-composite-next-hop scenario, when receiving a TTL expired packet, the device will transmit an ICMP error message in a MPLS header, but the route next hop for this ICMP error packet is discarded, so the one error message will be logged. [PR1194446](#)
- If RSVP link-protection optimize-timer is enabled, rpd memory might leak in "TED cross-connect" when a bypass LSP is being optimized. [PR1198775](#)
- If LDP neighbor relationship is over an unnumbered and flapping interface, the LDP will fail to advertise label binding. [PR1202071](#)
- With nsr and LDP export policy or l2-smart-policy configured, rpd on the backup Routing Engine might generate a core file when LDP is trying to delete a filtered label binding. [PR1211194](#)
- Due to an imperfect fix for a compatibility issue between 64-bit rpd and 32-bit client applications (such as "mpls ping", "monitor label-switched-path", and "monitor static-lsp") on 15.1F5-S3/15.1F6/14.2R7/15.1R4/16.1R1, the function of monitoring signaled or static LSP is broken on either 64-bit or 32-bit rpd. But the other 32-bit client applications (such as "mpls ping" ) are not impacted. [PR1213722](#)
- If the link/node failure that triggered a bypass persists for a long time, and there are LSPs that do not get globally repaired, multiple stale LSP entries are getting listed multiple times in the MPLS LSP. [PR1222179](#)
- This issue occurs in a multi-instance RSVP scenario with MPLS supported in the VRF routing-instance but the Connections protocol is not inside the VRF routing instance. When you are adding any interface under MPLS inside the VRF routing-instance, it should affect the Connections protocol inside the main instance. However, when you add the CE facing interface under MPLS in the VRF instance the Patricia with CCC information was deleted (because the CCC information was not inside the VRF instance). To resolve this issue, you would add a check that before acting on the Connections protocol, a check for whether the instance passed was master instance or not would occur. If it was not the master instance, the functionality related to CCC is not triggered. [PR1222570](#)

### **Network Management and Monitoring**

- Traps are sent as AgentX messages type (AGENTX\_MSG\_NOTIFY) from the subagent to the master agent. The subagent expects a response in form of an acknowledgement from SNMPD after sending these AGENTX\_MSG\_NOTIFY messages upstream. If an ACK is not received from snmpd within 1 second (current timeout value) the subagent will resend the trap. After router reboot or GRES, a lot of upstream communication is

triggered from subagent to snmpd (traps/mib registration messages). At that time, snmpd might not be able to send the downstream acknowledgement within a 1-second period. This might trigger the sub-agent to resend the trap, which will be seen as a duplicate trap on the NMS. To fix this issue, the timeout value has been increased from 1 second to 5 seconds in the sub-agent. [PR1164848](#)

- A trailing newline was erroneously added to the `$$message` variable. This had undesirable effects for some use cases when using the 'event-options policy `<>`' then `execute-commands commands <>` stanza. The fix escapes any newline characters and mitigates the issue. [PR1200820](#)
- Duplicated entries and error while loading MIBs on ManageEngine MIB Browser are fixed for the following MIB files: `jnx-chas-defines.mib`, `jnx-gen-set.mib`, `jnx-ifotn.mib`, and `jnx-optics.mib`. [PR1216567](#)

### ***Platform and Infrastructure***

- The **`show interfaces mac-database mac-address <mac-addr> <intf-name>`** command does not display any MAC-specific traffic statistics data on Stout Line cards and VMX for MAC-learning enabled interfaces mapped to the inet family. [PR1012046](#)
- Once the Traffic Offload Engine thread is stalled because memory error at the lookup chip, all statistics collected from the interfaces hosted by the Packet Forwarding Engine are not updated anymore. [PR1051076](#)
- Under large-scale setup, VPLS MAC might not be aged-out from remote Packet Forwarding Engine when local-Packet Forwarding Engine is MPC3/MPC4/MPC3E/MPC4E. As a result, unknown-unicast frames flood will be seen on local Packet Forwarding Engine. [PR1099253](#)
- This PR prevents an issue in which you could end up with two "`<junos:comment>`" entries under the `[interfaces]` stanza. [PR1102086](#)
- In software versions that contain [PR 1136360](#)'s code changes on MX Series Virtual Chassis (MX-VC), when J-Flow is not configured and equal-cost multipath (ECMP) load-balanced routes occur, the linecards might stop forwarding packets after logging any of the following errors prior to possible linecard restart or offline: - PPE Thread Timeout Traps - PPE Sync XTXN Err Trap - Uninitialized EDMEM Read Error. - LUCHIP FATAL ERROR - `pio_read_u64()` failed . In software versions that do not contain [PR 1136360](#) solution, on MX-VC with "virtual-chassis locality-bias" configured, ECMP load-balancing is occurring in the Virtual Chassis system, multicast streams and flooded Layer 2 streams might be duplicated or lost. [PR1104096](#)
- CoS error messages might appear when a nonexistent path for a database file is configured for CoS. These messages do not affect any service and traffic. [PR1158127](#)
- This issue occurs on MX Series routers with MPC6E linecards. The MPC6 only has 2 PICs (PIC number 0/1). If you try to configure an si interface with PIC number beyond the range (PIC number 2) on the MPC6E, it might crash, and traffic forwarding might be affected. [PR1160367](#)
- In a CoS environment with shaping-rate configuration under the interface, if flapping occurs on that CoS interface, the shaping-rate function does not take effect. [PR1163147](#)

- The error message, **CHASSISD\_UNSUPPORTED\_FPC: FPC with I2C ID of 0x0** is not supported can be seen on MX2020 after one FPC is pulled out and the configuration related-interface cannot be committed. [PR1164512](#)
- In some scenarios, when multiple logical systems are configured on a single physical router, an ordering issue might occur while updating routing states across the logical systems, causing the kernel to crash. [PR1169505](#)
- If configure micro-bfd on aggregate interface, when using native-vlan and if native-vlan is configured on one of the logical interfaces, then ARP resolution is failing for that IFL. [PR1172229](#)
- On all Junos OS platforms, when using RADIUS server, after RADIUS request is successfully sent by Junos device, if the network goes down suddenly, then response sent by the RADIUS server is not received within timeout period. In this scenario, the RADIUS request will be sent again with invalid socket descriptor, which will lead to auditd (provides an intermediary for sending audit records to RADIUS and/or TACACS+ servers) crash. [PR1173018](#)
- On MX Series routers with MPCs, if one AE interface is configured with a shared-bandwidth-policer, when the member link of the AE interface is removed or added, the policer value might be wrong. [PR1173704](#)
- The **show arp** command might not display complete results and reports "error: could not find interface entry for given index" due to some interfaces getting deleted when the show command is running. It is a timing issue and rarely happens. [PR1174150](#)
- This PR fixes an FD (file descriptor) leak problem in MGD process when netconf traceoptions are set. If <commit> rpc is executed via netconf session, there is an FD leak in the corresponding MGD pid. [PR1174696](#)
- When performing a unified ISSU (FRU upgrade) on EX9200-40T, EX9200-40F, EX9200-40F-M, EX9200-32XS, EX9200-2C-8XS, and EX9200-4QS line cards, an issue occurs with the buffer size in the line cards. As a result, the unified ISSU cannot be performed on EX9200 switches with these line cards. [PR1175240](#)
- On MX2020/2010, chassisd file rotation on commit check will cause the trace file to be stuck and no other operational chassisd events will be logged until chassisd restart. [PR1177625](#)
- When graceful Routing Engine switchover (GRES) is configured, the ksyncd crashes on the backup Routing Engine (RE) if a VPN static route has a network address as a next hop. As a result, the backup Routing Engine is not ready for graceful switchover. [PR1179192](#)
- After GRES, if you commit on the new master Routing Engine during the configuration synchronization with the old master Routing Engine, the commit might fail. [PR1179324](#)
- Rarely in an IPv6 sampling environment, IPv6 routes flap and the Packet Forwarding Engine might crash. This is a corner case; it is hard to reproduce. [PR1179776](#)
- On MX2K, the 'commit full' operation, or committing configuration under 'system' stanza (such as root-authentication and fxp0 interfaces) can cause transient Fan check Major alarm and Fan full speed. The Fan Tray spins at full speed for a while, then goes back to normal with clearing the alarm. The Fan check alarm and corresponding snmp



trap are temporal, and they can be safely ignored. user@router> show chassis alarms  
 2 alarms currently active Alarm time Class Description 2016-05-17 19:49:57 JST Major  
 Fan Tray X Failure 2016-05-17 19:49:57 JST Major Fan Tray Y Failure user@router>  
 show chassis environment Class Item Status Measurement Fans Fan Tray X Fan 1  
 Check Fan Tray X Fan 2 Check Fan Tray X Fan 3 Check Fan Tray X Fan 4 Check Fan  
 Tray X Fan 5 Check Fan Tray X Fan 6 Check Fan Tray Y Fan 1 Check Fan Tray Y Fan 2  
 Check Fan Tray Y Fan 3 Check Fan Tray Y Fan 4 Check Fan Tray Y Fan 5 Check Fan  
 Tray Y Fan 6 Check When MPC9E is installed in MX2K, the Fans usually keep around  
 6K rpm, and the fan speed control is frequently done by the Junos software. In this  
 situation, when all daemons are re-evaluated (by commit full or config change under  
 system stanza), the software bug causes the fan status to be checked within quite  
 small period, then the Junos software recognizes that the fan is faulty because the fan  
 speed has not reached the target speed yet when the fan status is checked within the  
 small period. After the fan alarm is detected, the fans are expected to start working  
 with full speed to cool the system components. The fan status check logic is fixed by  
 this PR. The fan status is checked after the fan speed is stabilized, hence we do not  
 see this transient fan alarm. [PR1185304](#)

- Auto-completion does not work for "show configuration" operational command for users with low privilege levels when allow/deny commands are configured. [PR1187130](#)
- From Junos OS release 15.1F2/16.1R1 and above, for MX Series routers with MPCs with inline-services enabled, the MPC crash might be observed during increasing the size of flow table by changing the "inline-services flow-table-size". [PR1188340](#)
- IPFIX Inline Sampling for IPv4 traffic might show incorrect flow distribution across ECMP next hops with **nexthop-learning enable** command configured. [PR1188545](#)
- When twamp server is configured with "routing-instance-list", and if the target-address configured on the twamp client is an si- interface address on twamp server, the twamp server might not work. [PR1189194](#)
- On MX240/MX480/MX960/MX2010/MX2020 Series routers with MPC used and configure firewall filter with scaling terms, when MPC boots (insert or reboot) or when any changes are made to that filter, the MPC might not come up. [PR1189669](#)
- In a very rare scenario, during TAC accounting configuration change, auditd daemon crashes due to a race condition between auditd and its sigalarm handler. [PR1191527](#)
- Insertion of an offlined MPC6E into the MX2K chassis can cause the FPC Temp sensor to detect transient "WARM TEMP" condition, and the chassis FAN in the same zone goes to high speed. \*\*\* messages \*\*\* Jul 12 18:10:17.698 MX2K-re0 chassisd[xxxx]: CHASSISD\_SNMP\_TRAP7: SNMP trap generated: FRU insertion (jnxFruContentsIndex 7, jnxFruL1Index 3, jnxFruL2Index 0, jnxFruL3Index 0, jnxFruName FPC: MPC6E 3D @ 2/\*/\*, jnxFruType 3, jnxFruSlot 2) MX2K-re0> show chassis zones |refresh 2  
 --- (refreshed at 2016-07-12 18:10:18 JST) --- ZONE 0 Status Driving FRU FPC 2  
 Temperature 63 degrees C / 145 degrees F Condition WARM TEMP  
 <----- Warm temp is detected Num Fans Missing 0 Num Fans Failed  
 0 Fan Duty Cycle 27 ZONE 1 Status Driving FRU SFB 5 SFB-XF2-Zone1 Temperature  
 59 degrees C / 138 degrees F Condition OK Num Fans Missing 0 Num Fans Failed 0  
 Fan Duty Cycle 27 --- (refreshed at 2016-07-12 18:10:20 JST) --- ZONE 0 Status Driving  
 FRU FPC 2 Temperature 63 degrees C / 145 degrees F Condition WARM TEMP



<----- Warm temp is detected Num Fans Missing 0 Num Fans Failed 0 Fan Duty Cycle 27 ZONE 1 Status Driving FRU SFB 5 SFB-XF2-Zone1 Temperature 59 degrees C / 138 degrees F Condition OK Num Fans Missing 0 Num Fans Failed 0 Fan Duty Cycle 27 --- (refreshed at 2016-07-12 18:10:22 JST)--- ZONE 0 Status Driving FRU FPC 2 Temperature 63 degrees C / 145 degrees F Condition OK Num Fans Missing 0 Num Fans Failed 0 Fan Duty Cycle 27 ZONE 1 Status Driving FRU SFB 5 SFB-XF2-Zone1 Temperature 59 degrees C / 138 degrees F Condition OK Num Fans Missing 0 Num Fans Failed 0 Fan Duty Cycle 27 Jul 12 18:10:27.489 MX2K-re0 chassisd[xxxx]: Fan Tray 0: zone 0 fan\_speed current 27% target 50% raising ratio 0.80 (linear) FPC 2 temp 72 last 72 WTC 55 WT 60 high limit 75 i2c\_ratio 0.80 Jul 12 18:10:27.490 MX2K-re0 chassisd[xxxx]: Fan Tray 0: set fan\_speed to 50% cfg\_speed 50% (linear) Jul 12 18:10:27.492 MX2K-re0 chassisd[xxxx]: Fan Tray 1: zone 0 fan\_speed current 27% target 50% raising ratio 0.80 (linear) FPC 2 temp 72 last 72 WTC 55 WT 60 high limit 75 i2c\_ratio 0.80 Jul 12 18:10:27.492 MX2K-re0 chassisd[xxxx]: Fan Tray 1: set fan\_speed to 50% cfg\_speed 50% (linear) Jul 12 18:10:47.517 MX2K-re0 chassisd[xxxx]: Fan Tray 0: zone 0 fan\_speed current 50% target 27% falling ratio 0.00 (linear) SFB 2 SFB-XF0-Zone0 temp 63 last 63 WTC 70 WT 75 high limit 90 i2c\_ratio -0.60 Jul 12 18:10:47.517 MX2K-re0 chassisd[xxxx]: Fan Tray 0: set fan\_speed to 27% cfg\_speed 27% (linear) Jul 12 18:10:47.519 MX2K-re0 chassisd[xxxx]: Fan Tray 1: zone 0 fan\_speed current 50% target 27% falling ratio 0.00 (linear) SFB 2 SFB-XF0-Zone0 temp 63 last 63 WTC 70 WT 75 high limit 90 i2c\_ratio -0.60 Jul 12 18:10:47.520 MX2K-re0 chassisd[xxxx]: Fan Tray 1: set fan\_speed to 27% cfg\_speed 27% (linear) [PR1193273](#)

- A rare VMCORE can occur because the process limit been breached by the creation of too many RSHD children processes. [PR1193792](#)
- When a configuration results in more than 63 bytes being placed in the encapsulation area, class-of-service traffic rate-limit is not correct for the small packets are (lower or equal to 77 bytes) or some packet is dropped. [PR1199853](#)
- On Junos OS platforms with command "delta-export" enabled, the delta-export database might not get correctly reinitialized upon one of the following conditions: 1. delta-export is enabled for first time (delta-export is enabled in just this commit), 2. load override (delta-export is enabled in the config), and 3. commit full (delta-export is enabled in the config). Data mismatches occur in further commits. As a result, the configuration on backup Routing Engine will be corrupted. [PR1199895](#)
- When checking default configurations about groups junos-defaults, no information is shown. [PR1201380](#)
- After system start up or after PSM reset you might see "PSM INP1 circuit Failure" error message. [PR1203005](#)
- When a Netconf ' <get-route-information>' RPC is executed for all routes via ssh transport session and the session is terminated before all the route information is retrieved, the MGD process and rpd daemon will cause high CPU utilization for an extended period of time. Examples of issues caused by this high CPU utilization for an extended period follows: BGP neighbors holddown timer expires and become ACTIVE OSPF adjacencies reset during database exchange OSPF LSA retransmissions events on neighboring nodes due to missing ACKs LDP sessions time out .Nondistributed BFD sessions are reset due to missing keepalives. [PR1203612](#)

- From Junos OS 15.1F2/14.2R4, validating configuration fails if commit scripts are used during software upgrade. [PR1204881](#)
- If inline J-Flow is configured in scaled scenarios, the sampler route database takes a long time to converge. [PR1206061](#)
- When "commit confirmed" is used after performing some changes, and an empty commit is performed to confirm the changes, the change-related processes will be notified again, which is unnecessary and might cause session/protocol flap. [PR1208230](#)
- On MX Series platforms with both DPC/E and MPC installed, when DPC/E detects a remote destination error towards a an MPC Packet Forwarding Engine, unexpected fabric drops occur. [PR1214461](#)
- In large-scale configurations or environments with high rates of churn, the FPC ASIC memory will become "fragmented" over time. It is possible in an extreme case that memory of a particular size will become exhausted and due to the fragmentation, the available memory will not fulfill the pending allocation. [PR1216300](#)
- On MX2K, MIC output is seen when there is no MIC in MPC under "show chassis hardware detail". Steps to reproduce the issue: 1. Offline MPC. 2. Physically remove MPC. 3. Physically remove MIC from the MPC. 4. Reinsert MPC. 5. Online MPC user@router> show chassis hardware detail | find fpc FPC 0 REV 68 750-044130 ABDxxx79 MPC6E 3D CPU REV 12 711-045719 ABDxxx35 RMPC PMB MIC 0 REV 14 750-049457 ABCxxx22 2X100GE CFP2 OTN >>>>>>> No MIC inside MIC 1 REV 26 750-046532 ABCxxx53 24X10GE SFPP >>>>>>>> No MIC inside XLM 0 REV 13 711-046638 ABDxxx59 MPC6E XL XLM 1 REV 13 711-046638 ABDxxx87 MPC6E XL [PR1216413](#)
- Trio-based linecards might crash after firewall filter configuration change is committed. [PR1220185](#)
- When any MPC line card is offlined, it goes offline via all offline flows and the connection is cleaned, but in the end of the offline flow, somehow it delays powering off the line card. The chasd powers off the MPC via and I2cs write the respective power registers, but the hardware is not really powering off. As a consequence, since the MPC is still powered-on but the connection is down, it will try to reconnect, then start to come up automatically within 10 seconds. [PR1222071](#)
- NTP peers fail to synchronize in symmetric active mode when there is significant downtime of one peer (for example, due to power maintenance, hardware or software upgrades). [PR1222544](#)
- Firewall filter index mapping become incorrect after Routing Engine switchover, because the contents of "/var/etc/filters/filter-define.conf" are incorrectly changed after Routing Engine switchover. [PR1230954](#)
- Incoming interface index cannot be used as a load- balancing input factor under family multiservice if the traffic payload is non-Ethernet frame. [PR1232943](#)
- Using < (lower-than-sign) as the first character of an item name under the groups hierarchy causes unexpected behavior and in the case of mpls LSPs causes the doubling of the configured LSPs. The '<' character is special when used under the "edit groups" stanza because there it is used to mark a "wildcard" name. Basically under "edit groups" you cannot use '<' at the beginning of any name. If you do use '<' at the beginning of

the name then '>' will be added automatically at the end of the name. In any other configuration hierarchy expect for "edit groups" you can use '<' and it will be considered literally. The documentation for configuration groups wildcards is available here:

[https://www.juniper.net/documentation/en\\_US/junos/topics/concept/junos-wildcard-notation-in-config-groups-use.html](https://www.juniper.net/documentation/en_US/junos/topics/concept/junos-wildcard-notation-in-config-groups-use.html) and it states that: "Wildcarding in configuration groups follows the same rules, but any term using a wildcard pattern must be enclosed in angle brackets <pattern> to differentiate it from other wildcarding in the configuration file." This PR introduces an extra check and warning to prevent the wrong use of the '<' character under the "edit groups" and thus prevent any unexpected behavior: [edit groups GRP\_MPLS\_LSPS protocols mpls] user@router# load merge terminal relative [Type ^D at a new line to end input] label-switched-path <TEST>\_LSP\_to\_3333\_no1 { to 3.3.3.3; } label-switched-path <TEST>\_LSP\_to\_3333\_no2 { to 3.3.3.3; } terminal:1:(42) Invalid groups wildcard notation, missing closing '>': <TEST>\_LSP\_to\_3333\_no1 [edit groups GRP\_MPLS\_LSPS protocols mpls label-switched-path] 'label-switched-path <TEST>\_LSP\_to\_3333\_no1 {' Invalid groups wildcard notation, missing closing '>' terminal:3:(1) error recovery ignores input until this point: } [edit groups GRP\_MPLS\_LSPS protocols mpls label-switched-path] '}' error recovery ignores input until this point ^D terminal:6:(1) error recovery ignores input until this point: } [edit groups GRP\_MPLS\_LSPS protocols mpls label-switched-path] '}' error recovery ignores input until this point load complete (3 errors) [edit groups GRP\_MPLS\_LSPS protocols mpls] user@router# show [PR1156024](#)

- If the user enters configuration mode with "configure exclusive" command after configuration is automatic rollback due to committing un-confirmed, the user still can make configuration changes with "replace pattern" command. The subsequent commit fails with "error: access has been revoked". After exit configuration mode, the user fails to enter configuration mode using "configure exclusive" with "error: configuration database modified". [PR1210942](#)
- When persist-groups-inheritance is configured and you issue a rollback, it will be seen that the configuration is not propagated properly after a commit. [PR1214743](#)

### VPNs

- In a multihomed source topology in NG-MVPN (applicable to both inter-AS and intra-AS scenarios), there are two problems. The first problem is Multicast (S, G) signaling does not follow RPF. When the routing table (mvpninstancename.inet0) has two routes, due to the policy configuration, the best route to the source is via the MPLS core, but Multicast (S, G) PIM join and NG-MVPN Type 7 both point to inactive route via the local BGP peer. The second problem is when "clear pim join instance NG" is entered, the multicast forwarding entries are wiped out. [PR1099720](#)
- After a GRES with NSR enabled, in NG-MVPN scenario, on the new backup Routing Engine RPD is consuming more than 90% CPU. This issue happens rarely and it is not reproducible. [PR1189623](#)
- In a BGP VPLS environment, sometimes you receive routes from BGP with invalid next-hop related information. In such scenarios, VPLS should treat them as bad routes and not send them to rpd infra for route resolution. Due to a software defect, the bad routes are passed to the route resolver, which might lead to rpd process crash. The

routing protocols are impacted and traffic disruption will be seen due to loss of routing information. [PR1192963](#)

- With MVPN and NSR enabled, high CPU on backup Routing Engine might be seen. MVPN on backup Routing Engine is re-queuing c-mcast events for flows because it is unable to find phantom routes from master RE. However, because routes is not reaching from the master RE, the backup Routing Engine keeps trying, causing high CPU utilization triggered by rpd processing. [PR1200867](#)
- Processing L2CKT/L2VPN/VPLS configuration might lead to a small memory leak in the routing protocol process (rpd). If the L2CKT/L2VPN/VPLS configuration is committed without any error, then there are 84 bytes in rpd memory that are not freed up and eventually exhaust of memory for rpd. [PR1220363](#)
- In a next-generation MVPN scenario with asm-override-ssm knob for source-specific multicast (SSM) group, if you issue "clear pim join" command on the source PE device, downstream interfaces get pruned causing the multicast flow to stop. [PR1232623](#)
- PR 1238807 introduces the ability to attach the connector ID attribute to all VPNv4 routes of a Draft Rosen based Multicast VPN (MVPN) instance. It is enabled with the "connector-id-advertise" keyword in the context of [routing-instances INSTANCE] and generates a type 1 connector ID attribute which preserves the BGP protocol next-hop (typically the PE router's loopback address) when the VPNv4 route is exchanged between autonomous system boundary routers (ASBR) of different autonomous systems (AS). This is important in scenarios where a MVPN stretches multiple AS and the AS-interop option B is used. [PR1238807](#)
- L2circuit does not switch from primary to backup and vice versa based on the APS status change. [PR1239381](#)
- With the NSR enabled and the l2circuit configured, rpd crash might be observed on the backup Routing Engine when you change the l2circuit neighbor and then commit the changes. [PR1241801](#)

#### **Resolved Issues: 15.1F6**

---

- [Class of Service \(CoS\) on page 93](#)
- [Forwarding and Sampling on page 93](#)
- [General Routing on page 95](#)
- [High Availability \(HA\) and Resiliency on page 107](#)
- [Infrastructure on page 108](#)
- [Interfaces and Chassis on page 108](#)
- [Layer 2 Features on page 110](#)
- [MPLS on page 111](#)
- [Network Management and Monitoring on page 114](#)
- [Platform and Infrastructure on page 114](#)
- [Routing Policy and Firewall Filters on page 120](#)
- [Routing Protocols on page 120](#)

- [Services Applications on page 124](#)
- [Subscriber Management and Services on page 125](#)
- [User Interface and Configuration on page 125](#)
- [VPNs on page 126](#)

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

- This PR does optimization in AE SNMP handling. If all the links in an AE bundle go down, then any COS SNMP query for this AE IFD/IFL will return cached values. [PR1140440](#)
- On MX104 platform, when applying the "rate-limit" and the "buffer-size" on the logical tunnel (lt-) interface on the missing MIC (not inserted on MPC), commit failure with error message would occur. As a workaround, this issue can be avoided by applying the "rate-limit and "buffer-size" on inserted MIC, then commit. [PR1142182](#)
- When customers delete an IFL from an interface-set that has CoS applied to it and activate CoS profile directly on that IFL in one single commit, commit fails with an error. Commit goes through if they do it one by one, delete IFL from interface set, commit and then activate CoS on that IFL, commit. [PR1169272](#)

### ***Forwarding and Sampling***

- Configuration statement "interface-mac-limit" might be set to default value when activating "mac-table-size" on a VPLS routing instance. Restarting l2ald, reapplying the "interface-mac-limit", or changing to another value (set interface ge-3/1/0.0 interface-mac-limit 510) fixes the issue. user@router> show vpls statistics | match count Current MAC count: 0 (Limit 1024) <<<<<<<<< set to default value 1024 instead of the value set by interface-mac-limit. [PR1025503](#)
- On MX Series platforms with MX-FPC/DPC, M7/10i with Enhance-FEB, M120, M320 with E3-FPC, when there are large-sized IPv6 firewall filters (for example, use prefix lists with 64k prefixes each) enabled, commit/commit check would fail and dfwd process would crash after configuration commit/commit check. There is no operational impact. [PR1120633](#)
- On MX80 and MX104 platforms, applying a firewall filter with an MX Series specific match condition will raise the following warning message: "Filter <filter\_name> is Trio specific; will not get installed on DPCs for interface <interface\_name>". This warning message is needed for the other modular-type MX Series platforms since they can have DPC and MPC mixed. But the message is not needed for MX80 and MX104 platforms since they only have the MX Series-based Packet Forwarding Engine. Although the warning message indicates that the relevant firewall filter is not installed, the firewall filter is correctly installed into the Packet Forwarding Engine. Thus, the user can ignore the message in case it is logged on MX80 and MX104 platforms. [PR1138220](#)
- The error message "pfed: rtplib: ERROR received async message with no handler: 4" might be seen when performing various operations on router (for example, when clearing BGP neighbors, or when doing BFD config). It means some messages sent by rtplib will not be received at pfed. (Please note, rtplib broadcasts events related to interface add, delete, modified etc. The pfe clients registered itself with rtplib library to listen for these events. In this particular scenario, rtplib was trying to call while pfed

was in middle of registration process, hence the error message. Messages may or may not be important for pfed). [PR1142836](#)

- For Junos OS Release 14.1R1 and later, when a broadcast packet is sent in a scenario of Integrated routing and bridging (IRB) over Virtual Tunnel End Point (VTEP) over IRB, the packet is getting dropped in kernel as it was looping due to a software issue. The error log message "if\_pfe\_vtep\_ttp\_output: if\_pfe\_ttp\_output failed with error 50" is observed when issue occurs. [PR1145358](#)
- On MX Series-based platforms, in race condition, when using the policer that has the configuration statement "bandwidth-percent" configured (e.g., set firewall policer XXX if-exceeding bandwidth-percent 80), if the logical interface (IFL) bandwidth change and the filter bind message arrive at the Packet Forwarding Engine out of order (e.g., when changing the bandwidth of the IFL or rebooting the FPC), the "bandwidth-percent" policer might end up using physical interface (IFD) bandwidth for "bandwidth-percent" computation. [PR1154034](#)
- When using MX Series-only features (gre decapsulate or payload protocol in IPv6), a change of policers or counters to an existing firewall filter using physical-interface-filter or interface-specific configuration statements will not be correctly detected by MIB2D. [PR1157043](#)
- This issue will be seen only when there are huge number of routes having different BGP NHs pointing to the same AS. Depending on the number of routes pointing to AS paths and also the difference in BGP NHs in the routes can shoot up the SRRD CPU consumption. In the real network this issue might not be seen often, as the number of AS paths will be huge and the routes referring these AS paths will be usually distributed among the AS paths. Even if the routes are pointing to the same AS, the impact would be lesser than the one seen in this PR. [PR1170656](#)
- When polling SNMP counters for MX Series-only firewall filters, MIB2D\_RTSLIB\_READ\_FAILURE cosmetic error messages might get reported in syslog. [PR1173057](#)
- Even if the packets do not match firewall filter conditions, wildcard mask firewall filter might match any packets. << Sample config >>
 

```

----- set firewall family inet filter TEST-filter
term TEST1 from destination-address 0.0.0.255/0.0.0.255 <<<<< set firewall family
inet filter TEST-filter term TEST1 then count TEST1 set firewall family inet filter
TEST-filter term TEST1 then discard set firewall family inet filter TEST-filter term
TEST2 then accept ----- This is discard filter
for /24 prefix broadcast address. However it might discard other packets. PR1175782

```
- This is a cosmetic issue. During sampling with jflow version 9, bfd packets from MPLS-TP were shown like as ip packets in "show services accounting aggregation template template-name XXX" command. (Actually, bfd packets info is not sampled by jflow.) << example >>
 

```

*****
lab@router-re0> show services accounting aggregation template template-name
mpls Src Dst Port/ Port/ Top MPLS MPLS MPLS Source Destination ICMP ICMP Label
Label 1 Label 2 Label 3 Address Address Type Code Proto TOS Address 299776 13 0
0.0.0.16 0.1.134.160 0 0 0 100.100.100.3 <<<<< bfd packet 299776 13 0 0.0.0.17

```

```

0.1.134.160 0 0 0 100.100.100.3 <<<< bfd packet 299776 16 0 10.0.0.1 40.0.0.2 8
0 1 0 100.100.100.3 <<<< ping 299792 16 0 40.0.0.2 10.0.0.1 0 0 1 0 100.100.100.1
<<<< ping 299776 16 0 40.0.0.2 10.0.0.1 0 0 1 0 100.100.100.3 <<<< ping
***** <
sample topology >>
*****
MPLS-TP(OAM, BFD) <-----> 10.0.0.1 40.0.0.2 sampling
[CE1]-----[PE1]-----[DUT]-----[PE2]-----[PE2] || [collector]
*****
PR1177876

```

- In Junos OS Release 15.1F5, family vpls filter applied to ae-interface is not working. [PR1178743](#)
- SRRD(Sampling Route-Record Daemon) process doesn't delete routes when the DELETE is received from rpd in few configuration cases. This results in build-up of memory in SRRD daemon and once SRRD reaches the limit, it crashes and restarts itself. This happens only when one certain family is not configured on all of the FPC clients (e.g. FPC with inline jflow enabled or PIC with PIC-based sampling enabled is one client). For example, only IPv4 family is configured in all the clients and, IPv6 and MPLS families are not configured for sampling in any of the clients. [PR1180158](#)

### General Routing

- On an MX Series Virtual Chassis platform, when we restart one or both of the standby Routing Engines, the log message "ksyncd\_select\_control\_plane\_proto: rhost\_sysctlbyname\_get: No such file or directory" might be observed as the ksyncd daemon attempts to select a communication protocol (UDP/TCP). After several tries, it will fall back to TCP and proceed as normal. [PR945925](#)
- An inconsistency between JUNIPER-VPN-MIB and MPLS-L3VPN-STD-MIB with the number of interfaces for a routing-instance has been identified. For example, with the following configuration: user@router-re0> show configuration routing-instances ri1 instance-type vrf; interface ge-2/0/8.10; interface lo0.10; route-distinguisher 65000:1; vrf-target target:65000:1; vrf-table-label; According to the MPLS-L3VPN-STD-MIB, there are two interfaces in this routing-instance: MPLS-L3VPN-STD-MIB :: mplsL3VpnVrfAssociatedInterfaces: OID: 1.3.6.1.2.1.10.166.11.1.2.2.1.8 Description: Total number of interfaces connected to this VRF (independent of ifOperStatus type). {master} user@router-re0> show snmp mib walk 1.3.6.1.2.1.10.166.11.1.2.2.1.8 mplsL3VpnVrfAssociatedInterfaces.3.114.105.49 = 2. However, according to JUNIPER-VPN-MIB there are three interfaces in this VRF: JUNIPER-VPN-MIB :: jnxVpnIfStatus OID: 1.3.6.1.4.1.2636.3.26.1.3.1.10 Description: Status of a monitored VPN interface. user@router-re0> show snmp mib walk 1.3.6.1.4.1.2636.3.26.1.3.1.10 jnxVpnIfStatus.2.3.114.105.49.733 = 5 jnxVpnIfStatus.2.3.114.105.49.754 = 5 jnxVpnIfStatus.2.3.114.105.49.774 = 5 The interfaces in the example are: {master} user@router-re0> show snmp mib walk 1.3.6.1.2.1.2.2.1.2 ifDescr.733 = ge-2/0/8.10 ifDescr.754 = lo0.10 ifDescr.774 = lsi.0 The fix for this issue adjusts this by removing the dynamic interface (in this case, lsi.0) from the interface list of JUNIPER-VPN-MIB. [PR1011763](#)



- On MX Series routers with MPC3E, MPC4E, MPC5E, MPC6E, Junos OS does not support short (sub-second) interface hold-time down configuration. So, a hidden configuration statement is introduced to ignore DFE tuning state during hold-down timer period. This configuration statement allows sub-second hold-down timer on MPC3E, MPC4E, MPC5E, MPC6E. set interfaces <intf name> hold-time up <U ms> down <D ms> alternative. The configuration statement does not work/support 'MPC5E 3D Q 2CGE+4XGE' and 'MIC6 2X100GE CFP2 OTN', and we recommend configuring hold-time down to be more than 3 seconds for these two cards. [PR1012365](#)
- The L2ald may crash after interface flap. [PR1015297](#)
- CoS scheduler names cannot be added or changed via service COAs. The schedulers can be added at subscriber login using client dynamic profiles. [PR1015616](#)
- No performance or functional impact. Can be safely ignored. "Ignore the PTP message (2) as this MPC does not support EEC" should be moved from notice to debug level. [PR1020161](#)
- MIC-3D-8OC3-2OC12-ATM Revision 22 or later is supported only by the following Junos OS releases: Junos OS Release 12.3 — 12.3R9 and later Junos OS Release 13.3 — 13.3R6 and later Junos OS Release 14.1 — 14.1R4 and later Junos OS Release 14.2 — 14.2R3 and later Junos OS Release 15.1 and later. [PR1036071](#)
- There is a remote loopback feature in 802.3ah standard, where one end can put the remote end into remote-loopback mode by sending an enable loopback control LFM PDU. In remote loopback, all incoming packets (except LFM packets) are sent back on wire as it is. Transmit or receive of LFM packets should not be affected when an interface is in remote loopback mode. On VMX platforms, when we configure the LFM remote-loopback we run into problem state. In problem state we will see that LFM packets sent from node which is in loopback state is not reaching the peer end, hence we will not see the remote entity information for the "run show oam ethernet link-fault-management" command on the peer router. [PR1046423](#)
- There are some configuration related functions in rpd and l2cpd that use special memory API called lite pools. These pools when reset were not freeing control information related to the pool and hence resulting in a leak. This is not a day one issue. This bug was introduced in 15.1 when we re-implemented LIBJTASK memory subsystem. This PR impacts all daemons using LIBJTASK (including rpd) on all platforms, provided memory lite pools are used by those daemons. [PR1071191](#)
- When flag is specified under ipsec-vpn traceoptions to trace IPsec operations, no message is logged to the specified trace file as expected. The issue impacts on debug capability only. [PR1073705](#)
- PCE-initiated LSPs are less preferred than locally configured LSPs. After this issue is fixed, PCE-initiated LSPs will have same preference as locally configured LSPs. [PR1075559](#)
- The Enhanced LAG feature is enabled in network-service enhanced-ip mode, but it is not supported in enhanced-ethernet mode. [PR1087982](#)
- When DHCP subscribers are terminated at specific routing-instances and the interface stack is IP demux over vlan-subinterface over AE interface, there might be a memory leak in the kernel AE iffamily when subscribers log in/log out. [PR1097824](#)



- In certain scenarios, Ksyncd might hog CPU while updating the if-state information to its peers depending on the volume of information being consumed by them. Following logs will be seen on the backup device : kernel: jlock hog timer expired: jlock acquired @ { 0xffffffff80b84ef8 0xffffffff80f08a06 0xffffffff8045a386 0xffffffff80b84dc0 0xffffffff8043ff5d 0xffffffff80439285 0xffffffff804395a2 0xffffffff80439634 0xffffffff8051c2fd 0xffffffff804f5181 }: thread 0xffff8016effe4f0, proc 0xffff8016ef73a30 (ksyncd), pid 6071, acquired msecs 2521 This might lead to CPU being not available for other critical operations involving kernel like TCP keep-alives and cause FPC disconnects. [PR1098534](#)
- With ECMP-FRR enabled, after rebooting the FPC which hoisting some ECMP links, the ECMP-FRR might not work. Clear any of BGP sessions (that is the part of ECMP) could help to clear this issue. [PR1101051](#)
- On MX Series routers where MS-MIC or MS-MPC is inserted, certain combinations of fragmented packets might lead to an MS-MIC or MS-MPC coredump. [PR1102367](#)
- During NSR Routing Engine switchover, there might be a control plane black window for inline BFD causing the BFD session to flap. This is a day-1 issue, and tuning the retrans timer would solve the problem. But since these timers have to meet RFC compliance, we cannot really do that. Today we have the retrans timer as 1000 milliseconds. The workaround would be to configure a higher retrans timer value. [PR1105980](#)
- On MX Series platforms, in rare conditions, if Packet Forwarding Engine sends wrong Packet Forwarding Engine id to chassisd as part of capability message, kernel might crash and some FPCs might be stuck in the present state, the traffic forwarding will be affected. This is a corner case, it is not reproduced consistently. [PR1108532](#)
- On MX240/480/960 Series routers with MS-DPC, customer is running BGP over IPSec. This BGP session has a BFD session tied to it. The BGP session is up but the BFD session remains in INIT state. The issue might be seen with any service configured with multi-hop BFD enabled. Traffic forwarding will not be affected. [PR1109660](#)
- Right now this fix is available from 14.2R6 and later. On 14.2R5 or earlier images, MSRPC gates once opened would never get deleted. From 14.2R6 and later, MSRPC gates are opened for 60 minutes no matter whether the expected packet hits the gate or not. After 60 minutes, gates are deleted by the timer. [PR1112520](#)
- Fixed problem with "egress pfe unspecified" increase when bind dhcp relay (or fpc restart caused ospf connection lose. Not able to ping its neighbor, arp table is fine, got egress pfe unspecified). [PR1114132](#)
- On MX Series routers containing multiple Packet Forwarding Engines such as MX240/MX480/MX960/MX2010/MX2020, with MPC3E/MPC4E/MPC5E/MPC6E cards, if the routers have GRE decap, then certain packet sizes coming via these line cards at very high rate can cause these line cards to exhibit a lockup, and one or more of their Packet Forwarding Engines corrupt traffic toward the router fabric. [PR1117665](#)
- During the LSP switchover, the hiwatermark might get set to an unexpectedly high value. The issue happens due to an incorrect reference point taken while calculating the Max avg BW in the last interval, and this results in an incorrect Highest Watermark BW in the autobandwidth stats. [PR1118573](#)

- On MX Series platform, in rare conditions, if removing or deactivating "member-interfaces" configured for an aggregated Multiservices (AMS) bundle (only officially supported on MS-MPC/MS-MIC), for example, using CLI command "deactivate interfaces ams0 load-balancing-options member-interface mams-7/1/0", all the MX Series-based FPCs and the MS-MPC/MS-MIC may crash. As a workaround, to avoid the issue, below is the recommended procedures to change AMS bundle size, 1. Offline member PICs. 2. Change AMS configuration. 3. Online member PICs. [PR1119092](#)
- On MS-MPC equipped MX Series platform, during the "three-way handshake" process, when receiving ACKs (e.g., after sending SYN and receiving SYN/ACK) with window size 0 (as reported, it is set to 0 by TCP client when using some proprietary protocol), the ACKs would be incorrectly dropped by the line card due to failure in TCP check. This issue could be avoided by preventing software from dropping packets that fail in the check, for example, by this CLI command, `re# set interfaces ms-3/0/0 services-options ignore-errors tcp`. [PR1120079](#)
- ANCP is not supported in this release. Attempts to use ANCP-related show commands will result in a timeout. [PR1121322](#)
- On MX Series platforms, the MS-MPC crash might occur. The exact trigger of the issue is unknown; normally, this issue might happen over long hours (e.g., within a week) of traffic run (e.g., running HTTP/HTTPS/DNS/RTSP/TFP/FTP traffic profile). [PR1124466](#)
- Right now this fix is available from Junos OS Release 14.2R6 and later. On Junos OS Release 14.2R5 or earlier images, SUN RPC gates once opened would never gets deleted. From Junos OS Release 14.2R6 and later, SUN RPC gates are opened for 60 minutes no matter whether the expected packet hits the gate or not. After 60 minutes, gates are deleted by the timer. [PR1125690](#)
- When Junos OS devices use the Link Layer Discovery Protocol (LLDP), the command "show lldp neighbor" displays the contents of PortID type, length, and value (TLV) received from the peer in the field 'Port Info', and it could be the neighbor's port identifier or port description. A Junos OS CLI configuration statement can select which "interface-name" or "SNMP ifIndex" to generate for the PortID TLV, so we do not have any problem as long as two Junos OS devices are connected for LLDP, but we might have an interoperability issue if another vendor device that can map the configured 'port description' in the PortID TLV is used. In this case, Junos OS displays the neighbor's PortDescription TLV in the Port info field, and if the peer sets the port description whose TLV length is longer than 33 bytes (included), Junos OS is not able to accept the LLDP packets and discards the packets as errors. The PortID TLV is given as : "the port id tlv length = port description field length + port id subtype(1B)". [PR1126680](#)
- If two redundant logical tunnel (rlt) sub-interfaces are configured in the same subnet and in the same routing-instance, a sub-interface will be down (this is expected), but if the sub-interface is removed from the routing-instance later, after disabling and enabling the rlt interface, a sub-interface might remain in the down state unless you remove the configuration of the rlt interface and then do a rollback. [PR1127200](#)
- rpd crash might be seen during deletion of address family on an interface while rpf check is configured. [PR1127856](#)
- In certain rare conditions, FPC VoQ will wedge which will drop packets on ingress Packet Forwarding Engine for MX Series router. Since the wedge is unable to be reproduced,

detection of wedge condition is introduced that alarm would be raised once the wedge condition is detected within 10 seconds. [PR1127958](#)

- On MX Series platforms with "subscriber-management" enabled, when a dynamic DHCPv4 subscriber is stacked over a static VLAN and the "route-suppression access-internal" knob is enabled, before the subscriber is established, it is possible for ARP process to first add a resolved route matching the subscriber's IP address. Then when the subscriber is established, the subscriber management process will change this route, but the change is not handled properly in the Packet Forwarding Engine. Due to this timing issue, the broadband network gateway (BNG) fails to forward transit packets to this subscriber. For example, the external DNS server's response packets might not be delivered to the voice subscriber interface, resulting in voice service outage. As a workaround, we can disable "route-suppression". [PR1128375](#)
- When using Point-to-Point Tunneling Protocol (PPTP) Application Layer Gateways (ALG) on MS-MPC/MS-MIC, if running scaled number of PPTP sessions control and data sessions (e.g., 1M sessions) for long hours (e.g., more than 8 hours), when the traffic is stopped, the "Bytes used" field of the output of CLI command "show services service-sets summary" will show a randomly large value due to memory issue. [PR1131605](#)
- In a situation where both mirrored interface and mirrored destination are on MPC card and mirror destination interface is a unilist next-hop (e.g., an ae interface), mirrored packets may get dropped. [PR1134523](#)
- The subscribers login rate could be degraded when IGMP/MLD is enabled on the dynamic demux interface. [PR1134558](#)
- Kernel crash might be seen due to integer wrapping around in case of 64-bit architecture. [PR1134578](#)
- From Junos OS release 14.1R4, 14.2R3, 15.1 and later, in the large scale environment (the scale is unknown), if restart NG-MPC ("request chassis fpc slot x restart or online"), it might not recover and remain offline. The traffic forwarding might be affected. [PR1135638](#)
- On MX Series platforms with non-Q-MPC (for example, MPC2-3D) or Q-MPC with enhanced-queueing off, when traffic has to egress on any one of the dynamic PPPoE (pp0), IP-DEMUX (demux0), and VLAN-DEMUX (demux0) IFLs, the queue mapping might be wrong. The traffic forwarding might be affected. [PR1135862](#)
- While bringing down subscribers, the system generates "Deinstantiate Service Failed permanently, daemon: cosd" error message. [PR1136083](#)
- On MX Series platforms with MIC3-3D-1X100GE-CFP, after in-service software upgrade (ISSU), the Junos upgrade is successful, but the 100GE port will be down, and the traffic forwarding will be affected. [PR1136269](#)
- In a IGMP oversubscriber environment with the configuration statement "remove-when-no-subscribers" configured, after performing graceful Routing Engine switchover, subscribers with multicast joins cannot re-login when the subscriber logs out before it sends an IGMP leave in the new master. [PR1136646](#)
- Insufficient time to allow an MPC5/MPC6 card to lock on the clocking source during FPC boot time might cause the Major Alarm raised due to "PLL Error." [PR1137577](#)

- On MS-MIC, TCP session Up/Down causes JSERVICES\_NAT\_\* and JSERVICES\_SESSION\_\* messages though severity level "none" are configured for services. [PR1137](#)
- For Point-to-Point Protocol over Ethernet over ATM (PPPoEoA) subscribers, during subscriber bringing down, the assigned IFL unit number is not correctly retrieved, so it can cause premature unit number exhaustion and thus fails to resolve &junos-interface-unit/&junos-interface-name variables. [PR1137723](#)
- In a multicast virtual private network (MVPN) scenario during route churn, the rpd process might crash due to inconsistency multicast next-hop between rpd and kernel. [PR1138366](#)
- In Junos OS Release 15.1F4, "show chassis environment" "Routing Engine 0 CPU" does not show Routing Engine CPU temperature as Junos OS Release 15.1F4 does not have the fix to make available the CPU temperature in Junos OS from HOST. "Routing Engine 0 CPU" instead shows the maximum of the Routing Engine inlet and exhaust sensors reading. [PR1140187](#)
- After removing a child link from AE bundle, the AE interface statistics in the SNMP MIB might show a spike. [PR1140533](#)
- FPC might restart while issuing "write coredump" from fpc shell. [PR1140870](#)
- From Junos OS Releases 14.1R4, 14.2R3, 15.1, and later, when a firewall filter is applied to NG-MPC, after system reboot, the Routing Engine might go into amnesiac mode. [PR1141101](#)
- Unending "mount request denied from 128.0.4.23 for /var/tmp/pics" messages are seen on the message log file. There is no functionality impact. Its just that it might overwhelm the Hard Disk with these messages. This would occur only with Service PIC being installed on one of the slots. [PR1141266](#)
- In subscriber management environment, on MX Series platforms, after login/logout static subscribers (e.g., by setting/deleting the interface), some of the static subscribers may be stuck in "Terminated" state. [PR1143205](#)
- When DHCP subscribers are brought up on the static interface IFL with interface-set, and this static interface IFL shares multiple DHCP stacks, it is possible that the interface-set does not get deleted when all DHCP subscriber are brought down on this static IFL. Unable to delete interface-set leads to commit denies on the dynamic profile involved. [PR1145450](#)
- Twice-NAT translation type does not work with the MS-MPC and MS-MIC service cards. The older MS-DPC cards does support his translation type. [PR1145690](#)
- On MX Series routers with "subscriber-management" enabled, the BBE subscriber management daemon (bbe-smgd) might crash on the backup Routing Engine when performing graceful Routing Engine switchover (GRES) during subscribers concurrent login/logout. [PR1147498](#)
- When a route in VRF has an indirect next hop, and the indirect next hop is pointing to a interface which is using un-numbered address, then the route in VRF table might be stuck in the KRT queue. [PR1147776](#)

- With a 100G CFP2 MIC installed in a MPC6E FPC, if the FPC fails to initialize the MIC, it is very likely that the FPC will get into boot loop. [PR1148325](#)
- Subscriber traffic in an LNS coming from the core network is not switched properly when the incoming interface is an irb interface. [PR1148533](#)
- On MX Series platforms, in multicast subscriber management environment (e.g., IGMP is configured for subscribers in dynamic profile), when nonstop active routing (NSR) is enabled, if the routing protocol process (rpd) is busy or there are hundreds of multicast groups are active (e.g., 250), missing multicast entries issue might be seen after performing Routing Engine switchover twice or more (i.e., first Routing Engine switchover works fine, and the issue may occur from the second switchover and onward). As a workaround, this issue could be avoided by issuing CLI command "restart smg-service" on backup Routing Engine after every switchover. [PR1149065](#)
- In EVPN environment, when CE MAC address alone gets changed for a MAC+IP entry, new MAC+IP entry is not getting reflected in the EVPN database and the old entry still exists on the PE router. [PR1149340](#)
- Commit error after attempting to delete all guaranteed rates on all traffic-control-profiles associated with demux0 [edit] lab@mx480-J12\_09# commit re0: [edit class-of-service interfaces] 'demux0' IFL excess rate not allowed on interface (demux0), please specify guaranteed rate on at least one IFL error: configuration check-out failed. [PR1150156](#)
- When using type 5 FPC on T4000 platform, traffic going out of the interface where "source-class-usage output" is configured will be dropped if the Source Class Usage (SCU) or Destination Class Usage (DCU) policy configuration is missing. This issue is caused by incomplete configuration, so to avoid the issue, please make the configuration complete (e.g., with "source-class-usage output" and SCU policy). [PR1151503](#)
- During deactivation of interfaces in a scaling setup, the Packet Forwarding Engine may reboot, resulting in traffic loss for a short period. [PR1151844](#)
- From Junos OS Release 14.2 with "exclude-hostname" configuration, hostname is not excluded from the messages before forwarding. This is a minor case, no other service impact. [PR1152254](#)
- Dynamic-tunnel interface bounces causing memory corruption leading to rpd crash. And the new rpd process once up, sync's up with the kernel, which may have information stored about the GRE tunnel ifl created by previous rpd process. The new rpd process using this information from the kernel leading to subsequent rpd crash being triggered. The following logs might be seen when this issue occurs: root@abc>show log messages| match "Address already in use" %DAEMON-3: Error creating dynamic logical interface from sub-unit 32792: Address already in use %DAEMON-3-RPD\_KRT\_Q\_RETRIES: kqp 0x49df00d0: op add queue low-add attempts 4010 ifd index 284, ifl unit 32792, family 2 instance id 0, state CreateIFL RPD\_KRT\_Q\_RETRIES: IFL IFF Update: Address already in use [PR1152912](#)
- OLD: set applications application my-ike-alg44 child-session-timeout 240 NEW: set applications application my-ike-alg44 child-inactivity-timeout 240 IKE ALG child sessions (ESP sessions) inactive timeout can be configured with this option. This option name is changed for better representation (the functionality is not changed). [PR1153045](#)

- Routers using inline layer 2 services may experience Packet Forwarding Engine wedge leading to fabric degradation and FPC restart. During issue state, the affected FPC will not be able to transmit and traffic will be fully blackholed. This problem is amplified by fragmented and out of order packets. This log entry may be seen during the error state: Host Loopback:HOST LOOPBACK WEDGE DETECTED IN PATH ID 0. [PR1153750](#)
- MPC7E/MPC8E/MPC9E control traffic is backing up and influenced during large-scale IS-IS convergence, cause LACP timeout and flapping. In addition, the entire system might be unstable and other protocols like IS-IS or LDP might also flap. [PR1154404](#)
- This feature is already available for physical MX Series router. Following is the link explaining about this feature. This PR provides same functionality on vMX also. [https://www.juniper.net/documentation/en\\_US/junos15.1/topics/reference/state-ment-hierarchy/system-services-resource-monitor.html](https://www.juniper.net/documentation/en_US/junos15.1/topics/reference/state-ment-hierarchy/system-services-resource-monitor.html) [PR1156184](#)
- CE in an EVPN setup which has no-mac-learning or is otherwise forwarding traffic upstream to MX's in an Active/Active EVPN configuration will see split horizon broken by the MX PE which has the MAC as DRC status. [PR1156187](#)
- Given an active BGP multipath route with 2+ ndirect-Next-Hops and another BGP route which can participate in protocol independent multipath with router-next-hop, rpd might crash if the interface on which first member of Indirect-Next-Hop resolves goes down. [PR1156811](#)
- A previous enhancement to strengthen the VC-Heartbeat message exchange resulted in rejecting messages at the crucial time of determining the health of the other VC member when all adjacency links fail. Validation of messages has been adjusted to remain strong when the VC is connected, but relaxed during the split conditions to prevent rejecting valid messages. [PR1157383](#)
- On MX Series platforms supporting MPC3E or MPC4E type MPC , the single-hop BFD session configured under VRF routing-instance can flap intermittently. The problem would be seen when the main-instance loopback firewall filter discards/rejects the BFD packets OR has term to accept only BFD packets from neighbors configured under main instance. In both scenarios, the BFD session packets coming on VRF routing-instance will be wrongly matched to main-instance loopback filter and gets discarded. With the fix of this PR, this situation is avoided and BFD session packets from VRF routing-instance will be matched with the correct VRF loopback filter (if configured). Note: In case there is no VRF loopback interface configured, then BFD packets are matched against main-instance loopback filter. [PR1157437](#)
- From Junos OS Release 13.2R1 and later, Packet Forwarding Engine interfaces on MX Series-based line cards might remain down after performing "request system reboot both-routing-engines " or "restart chassisd" several times. Rebooting the FPC might restore it. [PR1157987](#)
- On MX Series routers with MS-MICs and MS-MPCs, the Available addresses field in the output of the 'show services nat pool detail' command is always displayed as zero when destination NAT (dNAT) is configured. However, this field displays the correct number of addresses available for allocation when basic NAT or Network Address Port Translation (NAPT) is configured. [PR1158435](#)

- In PPPoEv6 scenario, the unsolicited Router Advertisement will be sent out before getting IPCPv6 ack. This behavior will impact PPPoEv6 connection rate. We can use "no-unsolicited-ra" configuration statement to suppress this message as a workaround. But in this case, this configuration statement does not work. The unsolicited Router Advertisement will still be sent out. [PR1158476](#)
- On MX Series platforms, when MPC experiences a FATAL error, it gets reported to the chassisd daemon. Based on the action that is defined for a FATAL error, the chassisd will take subsequent action for the FATAL error. By default, the action for FATAL error is to reset the MPC. When the MPC reports FATAL error, chassisd will send offline message and will power off the MPC upon the ACK reception. However, if MPC is in busy state for any reason, the ACK does not come in time and hence there would be a delay in bringing down the MPC. The fix ensures to bring down the MPC in time upon FATAL error. [PR1159742](#)
- In cases when the subscriber stacking is IPV6 over LNS, the IPV6 subscribers fails to come up with RPF check configured. DHC IPV6 subscriber over LNS comes up fine when RPF check configuration is disabled or removed. [PR1160370](#)
- Software OS thread on the the line card is doing a busy loop by reading the clock directly from hardware. Sometimes it seems the thread is getting wrong values from the hardware register and waiting forever in the busy loop. After the busy loop crosses a certain time period, the line card crashes and reboots. This is a rare condition. [PR1160452](#)
- On MX Series routers with enhanced queuing DPCs, there is a memory leak whenever doing SNMP walk to any of COS related OID's or issue the command "show interfaces interface-set queue <interface set name>". [PR1160642](#)
- The Router Lifetime field is set to 0 in the first Routing Advertisement sent from LNS back to PPPoE subscriber. [PR1160821](#)
- When FPC goes to terminated state (FPC down, restarts) ACI interface-set does not get deleted. After FPC comes online, further subscriber bring up on this ACI interface-set fails. [PR1161810](#)
- On Junos OS 15.1 and later, after Routing Engine switchover and both Routing Engine reboot, krt queue might get stuck. It's because: under this scenario, agentd creates its table before rpd reading tables. But after rpd restarting and rebuilding tables, it could not filter an agentd's table out. It might cause slow route convergence or traffic loss. This issue would disappear automatically in 30 minutes. [PR1162592](#)
- Changing VSTP configuration to MSTP configuration in one commit can result in service impact for some of the VLANS. Therefore it is advisable to do the configuration change in two steps. First deactivate VSTP and then activate MSTP. Same should be followed for MSTP to VSTP mode change. [PR1162661](#)
- During SIB yanking (pulling a SIB out without offline), it is possible that traffic may be dropped resulting in an overall reduction in traffic throughput. [PR1162977](#)
- The ICMP time exceeded error packet is not generated on an IPsec router on the decap side. The problem is fixed for MS-MPC/MIC and works fine if the session is there. There is no other way to return the time exceeded message over a tunnel. There is no plan to fix this for MS-DPC. [PR1163472](#)



- MQCHIP reports continuous "FI Cell underflow at the state stage" message and continuous fabric drops on ADPC ICHIP Packet Forwarding Engines after unified ISSU on MX Series with ADPC. [PR1163776](#)
- With MX Series platforms acting as TWAMP client and vMX platform acting as TWAMP server setup, we see probe packet loss at TWAMP server, i.e., on VMX with Junos OS Release 15.1F5. When the TWAMP target interface address is configured as a Media interface (-ge/-xe), probe packets are getting dropped at vMx because of ENDIAN conversion of UDP checksum (vMx is Little Endian and Mx is Big Endian platform) in the probe packet. This issue was seen earlier in Junos OS Release 15.1F4 but was resolved through PR1125516. However, due to some merge issue the fix got overwritten and this issue is resurfaced. Also, when the TWAMP target interface address is configured as si- interface, we again see probe-packet loss. but this time not because of UDP checksum error. Here, the issue appears because of some looping issue and packet after getting processed at LU (timestamped at LU) is not able to go out of the media interface. Sometimes enabling some debug logs at the Packet Forwarding Engine and changing TWAMP probe packet size resolves the issue (but not always).  
[PR1164093](#)
- From Junos OS Releases 14.1X51-D75, 15.1F4, and 15.2I1B, if deactivate interface, in rare condition, due to a software defect, NH is getting deleted, while there are still routes pointing to it, which leads to inconsistent states in Packet Forwarding Engine. The MPC might crash or traffic blackhole. [PR1164101](#)
- When using qsfp28 optics on 100G Gladiator PTX FPC card, the Tx laser disabled alarm is not on after disabling interface. This has been already fixed in 15.1F6. One of the function was set w/ incorrect value that was causing this problem when scanning Tx disable alarms . Laser bias current low alarm : On Laser bias current high warning : Off Laser bias current low warning : On Laser receiver power high alarm : Off Laser receiver power low alarm : Off Laser receiver power high warning : Off Laser receiver power low warning : Off Tx loss of signal functionality alarm : Off Rx loss of signal alarm : Off Tx laser disabled alarm : Off<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<<, Lane 3 Laser bias current : 0.000 mA Laser output power : 0.000 mW / - Inf dBm Laser receiver power : 1.091 mW / 0.38 dBm Laser bias current high alarm : Off Laser bias current low alarm : On Laser bias current high warning : Off Laser bias current low warning : On Laser receiver power high alarm : Off Laser receiver power low alarm : Off Laser receiver power high warning : Off Laser receiver power low warning : Off Tx loss of signal functionality alarm : Off Rx loss of signal alarm : Off Tx laser disabled alarm : Off<<<<<<<<<<<<<<<<<<<<<<<<<<<<<< [edit] lab@ptx5k# [PR1164163](#)
- The ability to configure a multicast group statically for a subscriber via a dynamic profile is not available in this release. Using the below statement, the subscriber can be enabled to receive multicast traffic for group 224.117.71.1 upon login: 'set dynamic-profiles <client profile> protocols igmp interface "\$junos-interface-name" static group 224.117.71.1' This support is not available and the subscriber needs to send a IGMP protocol JOIN message to receive multicast traffic. [PR1164323](#)
- With Junos OS Release 15.1 and later, on MS-MPC or MS-PIC, OSPF adjacency may fail to establish when there is no static route pointing to service PIC. [PR1164517](#)
- On a chassis populated with MPC5E/MPC6E and loaded with 15.1 image, MPC5E/MPC6E will drop traffic during FRU upgrade whereas other supported cards



will continue to undergo ISSU. This issue is caused by Shadow RAM be used after ISSU replay completes. [PR1165172](#)

- With IKEv1, MS-MPC packet drops on far-end after reboot of local MS-MPC. [PR1165787](#)
- Hardware has feature to shutoff the retimers if there is power in single zone . Ideally retimers from those zones should not be initialized .since this requires more testing planning to fix it in next release . [PR1168059](#)
- MS-MPC may crash when bridge domain is used. [PR1169508](#)
- When Estimated BER is concluded as just above the Transition BER, it should not be compared with the Thresholds, as the Estimated BER value is not exactly known at this point. [PR1169972](#)
- If a given demux VLAN hosts both dynamic IP demux subscribers as well as static IP demux interfaces, it is possible that the dynamic IP demux subscribers appear to bind successfully, but they can experience forwarding problems. In this scenario, the dynamic subscriber state is not fully established on the line card, resulting in traffic issues. [PR1170019](#)
- Adding keyword 'fast-filter-lookup' to existing filters of an input or output filter list may result in failure to pass traffic. To avoid this issue, the filter list should first be deactivated then the filters updated with a the keyword 'fast-filter-lookup; then the filter list activated. [PR1170286](#)
- If the "no-cell-share" configuration statement under the chassis stanza is activated on MPC3, MPC4, MPC5 or MPC6 cards, the Packet Forwarding Engine will only be able to forward about 62 Gbps versus ~130 Gbps, causing fabric queue drops. [PR1170805](#)
- `sctl_hwre_ngre_virtio_fixup()` is not SMP-safe and, as such, has the ability to corrupt jlock protected data, or even cause the Routing Engine to crash. [PR1172346](#)
- When upgrade Junos software on RE1, if at the time, RE1 is the "master Routing Engine", both Routing Engines will be in "backup" state. Resulting in loosing remote connectivity, and all interfaces. Only "console" access will be available at this time. [PR1172729](#)
- When upgrading or rebooting the router, the following logs might be seen in 15.1F5. There is no impact and they can be ignored. This is due to the fact that agentd is trying to read the forwarding class entries at system boot time too early, when they are not yet created. This has been fixed. <.> FILE SYSTEM CLEAN; SKIPPING CHECKS clean, 9762157 free (813 frags, 1220168 blocks, 0.0% fragmentation) tuneufs: soft updates remains unchanged as disabled chown: wheel: Invalid argument Creating initial configuration...agent for all the telemetry sensors: COSD\_CONF\_OPEN\_FAILURE: Unable to open: /var/etc/cosd.conf, using default CoS forwarding classes, do 'commit full' in cli to avoid this message agent for all the telemetry sensors: COSD\_CONF\_OPEN\_FAILURE: Unable to open: /var/etc/cosd.conf, using default CoS forwarding classes, do 'commit full' in CLI to avoid this message mgd: commit complete. [PR1173137](#)
- When using Periodic Packet Management process (PPMD, responsible for periodic transmission of packets on behalf of its various clients) related protocols (e.g. LFM, CFM, LACP, BFD, etc), during fabric or SIB online process, possibly, the client session (who establish adjacencies with PPMD to receive/send periodic packets on those

adjacencies, such as LFM, CFM, LACP, etc) of PPMD may flap due to CPU hog issue. [PR1174043](#)

- The fan speed logic does not operate correctly once PEM on MX104 platforms automatically shuts down due to over-temperature protection. The fan speed moves back to speed normal. It takes more time for PEM to cool down and come back online automatically with fan at normal speed. [PR1174528](#)
- When using MS-MPC or MS-MIC service cards, a single pool cannot be used in different service-sets. Separate pools with different names would then need to be used. Additionally, pools created automatically by a source-prefix or destination-prefix statement will not work if the same source-prefix or destination-prefix statement appears in a different service-set. [PR1175664](#)
- Inline-services flow-table resize without reboot might cause `jnh_write_lkup_inst_internal()` errors from being logged. [PR1176186](#)
- Storm control feature is not working on MX104 platform. In Packet Forwarding Engine, associated filters and vty commands are not visible as well. It works on other MX series platforms. [PR1176575](#)
- MACSEC not working on layer 3 interface on MX104 [PR1177630](#)
- destination-prefix-list support list added for NAT rule with twice-napt-44 translation. Customer will be able to define a prefix list and match it in the NAT rule while using twice-napt-44. [PR1177732](#)
- In a rare error scenario, `krt_q_entry` of flow route was freed without dequeuing it from queue. This has been fixed via software change. [PR1178633](#)
- In MX Series running a Junos OS Subscriber Management Build, with more than 300+ firewall filters configured, it was found that a subscriber failed to login due to NACK received from the system, stating the following error: "BBE\_DFW\_DYN\_PROF\_ERR\_STR session\_id=1784: Can't find filter template named test300. BBE\_DFW\_DYN\_PROF\_ERR\_CODE session\_id=1784: Error code 13: Filter template not found." While the firewall filter named "test300" was certainly configured under the firewall filter configuration stanza, it found that the BBE daemon could hold a count of 256 filters only. Filters above this count were not getting indexed into the internal filter table and hence system could not find the filter. [PR1178671](#)
- Changes are needed to support dedicated users for control and multicast traffic. This will avoid unicast traffic to be hashed to users doing ucode processing. On Junos OS side, this PR introduces new CLI command "set chassis fpc X performance-mode num-of-ucode-workers Y". [PR1178811](#)
- In EVPN A/S mode, IFL mark down programming at the sPacket Forwarding Engine on the BDF gets removed, causing traffic loops. [PR1179026](#)
- [EVPN] Active-Active IP4 L3 session with CE over IRB Flaps [PR1179105](#)
- On 10x10GE(LAN/WAN) SFPP PIC, when the port is configured with WAN PHY mode, the CoS configuration on the port will be incorrectly programmed and it might result in unexpected packet drop. [PR1179556](#)

- When a MPC has training failure on all planes, then other MPCs in the system are getting affected. The root cause is that MQ MPC are not deleting the streams of the MPCs which is causing the fabric wedge and effecting other MPCs. As a result FH is kicking in for other MPCs in the system. [PR1183230](#)
- In the CGNAT CLI show service alg conversations fails to display parent session status for ALG conversations. [PR1181140](#)
- On MX2010/2020 routers with SFB2 and empty fabric slots, a system defect that fetches wrong fabric info might cause MPC7E/8E/9E not being able to come online. [PR1182404](#)
- [EVPN] Active-Active IP4 L3 session with CE over IRB flaps. [PR1179105](#)
- ICMP pings destined to VIP address beyond 166 bytes are dropped as "my-mac check failed" [PR1186537](#)
- When LFM session is configured with timeout of 300ms or less, it might flap during another MPC's offline sequence. [PR1191546](#)

#### ***High Availability (HA) and Resiliency***

- On MX240/480/960/2010/2020 platforms with Junos OS Release 15.1R1 and later, in a high-scale scenario (e.g., there are 4 million routes or more), the connection between Routing Engine and the FPC(s) might flap after performing graceful Routing Engine switchover (GRES). The other symptoms are intermittent packet drops between the Routing Engine and FPC during regular operation without performing GRES and scaled scenario. [PR1146548](#)
- Unified ISSU between Junos OS Release 15.1F3 and earlier to Junos OS Release 15.1F4, and ISSU between releases from Junos OS Release 15.1F4 to Junos OS Release 15.1F5 will result in a core dump and could lead to PR1161491. The same might happen when the ISSU is done from 14.2R4/R5 to 15.1F5 only. This issue happens due to an inconsistency in port numbering between two port types in the releases. There could be other consequences due to this issue in the upgraded release that might hamper functionality on some types of ports only. [PR1161491](#)
- Right after all FPC complete their upgrade, the kernel (on the VC-Mm) closes its connection to ksyncd (on the VC-Bm) since it has received a message "invalid IPC type 20". This disconnect causes ksyncd to restart, it then cleans all kernel state in the VC-Bm and starts the replication process. This causes the timer for waiting for the VC to become GRES ready (after FPC upgrade) to expire and abort the ISSU.. [PR1163807](#)
- When configuring the "nonstop-routing" under one group and applying this group to routing-options configuration hierarchy, sometimes the NSR does not work. As a workaround, please configure the "nonstop-routing" directly under the routing instance hierarchy. [PR1168818](#)

### **Infrastructure**

- In scaling setup (in this case, there are 1000 VLANs, 1000 Bridge Domains, 120 IRB interfaces, 120 VRRP instances, BGP and IGP), if the routing protocols are deactivated and activated, there might be a chance that the pending route stats are not cleaned up, which will cause the stats infra to have stale pointers and lead to memory corruption in socket layers. The system might go to db prompt because of this. All the traffic going through the router will be dropped. [PR1146720](#)
- In Legacy, if the Routing Engine switchover on rpd crash configuration statement is enabled, the switchover is instigated in the Routing Engine kernel BEFORE the rpd core is created. In Occam, it is done AFTER. This creates an issue in scaled setups where the size of the rpd core, and therefore the time to create it, takes a lot longer. An Occam FreeBSD change has been made as part of this PR patch to adopt the same behavior as Legacy. [PR1164139](#)
- Support for fast Routing Engine mastership switchover when master Routing Engine crashes was broken in Occam releases up to and including 15.1R3. The patch for this PR addresses support for this feature in Occam in 15.1R4. [PR1167385](#)
- Issue is specific to Occam based images and is a very hard to induce. The issue occurs when virtual memory is low, and the FreeBSD OS tries to free-up memory by invoking the vm\_lowmem event. In a multi-core environment, multiple kernel threads could encounter the low virtual memory condition at the same time, and so the vm\_lowmem event could be invoked within the context of multiple kernel threads concurrently. Some of the protocols in the Junos networking stack register handlers against this event and two of these, clnp\_drain() & tcp\_drain(), were not SMP safe, which caused data corruption. clnp\_drain() & tcp\_drain() have now been made SMP safe; all other such handlers in the Junos networking stack were already SMP safe. [PR1182958](#)

### **Interfaces and Chassis**

- jnxBoxDescr is reworded for MXVC to replace the platform type with a more general representation that replaces the specific member platform type with "Virtual Chassis." Old virtual chassis text example: jnxBoxDescr.0 = member0 Juniper MX240 Internet Backbone Router. New virtual chassis text example: jnxBoxDescr.0 = member0 Juniper MX Virtual Chassis Internet Backbone Router. NOTE: The MIB design for jnxBoxAnatomy "top-level" chassis information works properly for a standalone chassis, but doesn't fully represent virtual chassis multi-member configurations because it is capable of providing information for only one physical chassis. (The remainder of the jnxBoxAnatomy MIB "containers" properly support the inventory of a multi-member configuration.) MX virtual chassis provides another MIB, jnxVirtualChassisMemberTable, to supply the equivalent "top-level" information. [PR1024660](#)
- MXVC-specific behavior for SNMP walk of jnxOperating\* containers was divergent from the physical MX Series. Returned to vergence. [PR1136414](#)
- %DAEMON-3-CHASSISD\_I2C\_WRITE\_ERROR: i2cs\_write\_reg: write error for group 8 at address 0x49, offset 32 %DAEMON-3-CHASSISD\_I2CS\_READBACK\_ERROR: Readback error from I2C slave for FPC 1 ([0x11, 0x42] -> 0x0) - The above errors represent transient communication issues between system components. -

In certain cases, these can be service impacting. - Enhancements have been made for better handling of such error conditions. [PR1139920](#)

- When micro Bidirectional Forwarding Detection (BFD) sessions are configured for link aggregation group (LAG), the device control process (DCD) acts as the client to the micro BFD session. In order to monitor the connection between client (DCD) and server(BFD), the client needs to exchange keepalive hello packets with the server. To send hello packets, DCD needs to move out of IDLE phase to CONFIG\_BFD phase, which is the reason for the following log messages: dcd.c:585 dcd\_new\_phase\_if\_idle() INFO: Current phase is IDLE, going to phase CONFIG\_BFD usage.c:75 dcd\_trace\_times() INFO: Phase Usage for IDLE : user 0.001 s, sys 0.000 s, wall 60.019 s dcd.c:717 dcd\_new\_phase() INFO: New phase is CONFIG\_BFD usage.c:75 dcd\_trace\_times() INFO: Phase Usage for CONFIG\_BFD : user 0.000 s, sys 0.000 s, wall 0.000 s dcd.c:717 dcd\_new\_phase() INFO: New phase is IDLE There is no functionality impact; however, these messages might flood the logs. As a workaround, we can filter out these messages from being written to the log file according to this KB article: <https://kb.juniper.net/InfoCenter/KB9382>. [PR1144093](#)
- In MX-VC or VRRP platforms running Junos OS Releases of 15.1 built before about February 2016, the following cosmetic warning message will be displayed upon commit: [edit] 'chassis' warning: WARNING: MPC reboot or chassis reboot is required to use MIC aware dynamic power management feature on already plugged-in MPCs. [PR1144295](#)
- During a VRRP configuration change involving IP address change and/or VRRP configuration change while retaining same group ID, a race condition might occur causing vrpd crash. [PR1145170](#)
- When using MX Series platforms as Layer 2 Tunnel Protocol (L2TP) access concentrator (LAC), if login/logout tunneled PPPoE subscribers over an extended period (e.g., login/logout 16K subscribers for 24- 48 hours), kernel crash may occur due to next-hop issue. [PR1150316](#)
- In affected releases, the following cosmetic alarms are seen after reseating the clocking cables: "2015-11-13 05:22:56 UTC Major CB 0 External-A LOS 2015-11-13 05:22:56 UTC Major CB 0 External-B LOS." [PR1152035](#)
- Remove MX series from sending LCD halt message. [PR1153219](#)
- The outbound PPPoE control packets sourced from PPPoE daemon, such as PADO, error PADS and PADT, are assigned to queue 0 instead of queue 3. [PR1154070](#)
- SONET interface on MIC-3D-IOC192-XFP does not count input error correctly. While hardware counts framing error, runts and giants but input error in 'show interface extensive' command reports runts, and giants only. [PR1154268](#)
- "monitor interface <if name>" will start ifmon process. In this time if telnet session to router is disconnected unconventionally, then ifmon process was not killed and it will take up 100% CPU utilization. The workaround is to terminate the stale ifmon process. [PR1162521](#)
- Internal timing for bringing FPCs online is extended for MX 2020/2010 systems to accommodate longer initialization times for fabric and FPCs. [PR1164147](#)

- CLI commit warning is replaced by syslog warning message when limited-ifl-scaling configuration statement is configured. Warning message text remains the same. [PR1165357](#)
- jpppd core at SessionDatabase::getAttribute() from Ppp::LinkInterfaceMsOper::getLowerInterfaceType() [PR1165543](#)
- During graceful-switchover, the packet forwarding engine will attempt to send the LFM packets to the new master routing-engine in order to refresh/create the adjacencies. If the connection to the new routing-engine is not yet ready, these packets will be delayed and LFM might flap. [PR1167760](#)
- If an interface configured with VRRP is removed from a routing-instance to global, or from global to a routing-instance, the IFLs of that interface will be deleted and recreated. In ideal case as the interface gets deleted, VRRP should move to bringup state, when the interface is created again, VRRP goes to previous state. After this VRRP should get VIP addition notification from kernel and update VRRP state and group id for VIP. However, in race conditions, VRRP might get VIP addition notification from kernel even before the interface creation event happens. If so, VRRP will never be able to update proper VRRP state and group id. So the VIP will reply for the ARP with an incorrect MAC ending with "00", while the correct MAC should end with the groups id configured. [PR1169808](#)
- When upgrading Junos OS software on RE1, if at the time, RE1 is the "master Routing Engine", both Routing Engines will be in "backup" state, resulting in losing remote connectivity, and all interfaces. Only "console" access will be available at this time. [PR1172729](#)
- Commit check may exit without providing correct error message and causing dcd exit. [PR1180426](#)

### **Layer 2 Features**

- There is a bug in code of handling the redistribution of PPM (periodic packet management) Transmit and Adjacency entries for LACP, when the Interface entry is in pending distribution state. This issue might cause pppmd crash after graceful Routing Engine switchover. [PR1116741](#)
- On GRES switch of mastership of Routing Engine via "request chassis routing-engine master switch", the dot1xd daemon will crash multiple times when 128K IFLs are configured in the MX960 chassis [PR1118475](#)
- In BGP-based VPLS scenarios, changing the configuration of a VPLS mesh group might cause rpd core. FPC reboot might also be seen during the rpd core. [PR1123155](#)
- On MX Series platforms, in DHCP subscriber management environment (the device is either used as local DHCP server or DHCP relay agent), if configuring the Aggregate Ethernet (AE) interface (e.g., change the "MTU" of AE) while there are subscribers on it, in race condition, the DHCP binding failure would occur on the AE. [PR1139394](#)
- In a VPLS scenario, when "\$junos-underlying-interface-unit" is configured in "dynamic-profiles" hierarchy, which is then implemented in a routing-instance, the upgrade/commit will fail with the following error message: "Parse of the dynamic

profile (<dynamic\_profile\_name>) for the interface: \$junos-interface-ifd-name and unit: \$junos-underlying-interface-unit failed!". [PR1147990](#)

- For router equipped with following line cards: T4000-FPC5-3D MX-MPC3E-3D MPC5E-40G10G MPC5EQ-40G10G MPC6E MX2K-MPC6E, If the router is working as VPLS PE, due to MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)
- From Junos OS Release 13.2R1 and later, the rpd process might crash when adding/deleting Virtual private LAN service (VPLS) neighbors in a single commit. For example, a primary neighbor is changed to become the backup neighbor. [PR1151497](#)
- In some cases where DHCP client devices are not fully protocol compliant they may become stuck trying to renew an address lease indefinitely. These devices exposed a defect in the DHCP Relay behavior when acting as a proxy for the server where a protocol NAK to restart the client was not properly created. As a result address resources could be locked on the relay, preventing their use until the offending client device was restarted. [PR1153837](#)
- In 15.1R3 with tomcat mode enabled, DHCP subscriber management with IRB interfaces is not reliable. It is possible that the DHCP bindings are unable to fully establish with IRB interfaces due to this reason. However, these bindings with same IRB interfaces should come up properly with tomcat disabled. [PR1155502](#)
- The "Node ID" information is not shown on MX Series platforms when traceoption flag "pdu" is configured to trace Ethernet ring protection switching (ERPS) PDU reception and transmission. [PR1157219](#)
- When an MX Series router is acting as DHCP relay to selectively process client traffic with any forward-only configuration, if a downstream device acts as a Layer 2 DHCP relay where it adds an OPTION-82 record but not a giaddr (Gateway Address field), and in addition, the downstream Layer 2 DHCP relay adds the option 82 record in a non-compliant (illegal) way by inserting the OPTION-82 record in front of other existing option records, bad packet format of the DHCP discover/request will send to the server. [PR1157800](#)

### **MPLS**

- In MPLS environment, the master Routing Engine might crash due to Mbuffer allocation failure and this crash will trigger an Routing Engine switchover, as a result backup Routing Engine will become active. The issue is unreproducible, and trigger condition is not clear. [PR979448](#)
- If a RSVP LSP has both primary and secondary standby path and link-protection enabled, a /32 bypass route is unhidden when the primary link goes down. This /32 route is supposed to be made hidden again when primary link comes back up. But in some cases, due to software defect, this /32 bypass route remains unhidden forever which causes some issues, for example, BFD session down due to better prefix received from bypass LSP. [PR1115895](#)
- During interoperation with Cisco device (e.g., CRS) belonging to different IGP area, if the P2MP LSP ping echo reply message from Cisco device is using interface address other than loopback/router-id as the source address, the reply message will be dropped



on Junos OS device. With the fix, Junos OS device will accept the packets and print them as 'uncorrelated responses'. [PR1117166](#)

- Due to some data structure changes of ipc messages in 64-bit rpd, some of 32-bit applications (e.g., lsping, lspmon) would not work normally when rpd is running in 64-bit mode. Depending on Junos OS version, some of CLI commands might not work as expected. [PR1125266](#)
- When an PLR is a non-Juniper router, Juniper ingress node might stay on the bypass tunnel and ignore the CSPF result. [PR1138252](#)
- When a link fails on an RSVP LSP which has link-protection or node-link-protection configured, the PLR (point of local repair) will initiate a bypass LSP and the RSVP LSP will be tunneled on this bypass LSP. However, if now the bypass LSP is brought down because there is a link failure on it, the PLR might only send out a session\_preempted PathErr message to the upstream node without sending a ResvTear message. Hence the ingress node does not receive a ResvTear message and the RSVP LSP is not immediately torn down. The RSVP LSP will remain UP for more than 2 minutes until the RSB (Resv state block) on the ingress's downstream node gets timed out and it sends a ResvTear message to the ingress. [PR1140177](#)
- There is no entropy label for LDP route in a scenario of LDP tunneling across a single-hop RSVP LSP with label 0 (explicit-null) used. As workaround, either remove LDP tunneling or RSVP explicit-null will resolve the issue. [PR1142357](#)
- This issue is related to inter-op between multivendors scenario. This fix will add sub-object RRO, which will help change of label during FRR active scenario. [PR1145627](#)
- MPLS TED might not select random links to calculate the ERO when OSPF is overloaded. Instead, only one or two interfaces will be used for all the configured LSPs originating from the router. [PR1147832](#)
- In LDP P2MP scenario with NSR, after performing multiple iterations of FPC reloads, protocol bounce, interface bounce, GRES, rpd restarts in random, in rare condition, the rpd process might crash, the routing protocols are impacted and traffic disruption will be seen due to loss of routing information. [PR1148404](#)
- With RSVP refresh reduction feature enabled (using RSVP aggregate messages), when changing the knob "no-load-balance-label-capability" to "load-balance-label-capability" on the egress router, the Entropy Label Capability (ELC) for the egress router would not being propagated towards the ingress. As a workaround, we can execute "clear rsvp session" on the ingress or wait until 3 refresh cycles (say 100s with default RSVP refresh config). [PR1150624](#)
- Static MPLS LSP using VT interface as a outgoing interface would not come up. [PR1151737](#)
- With NSR enabled and LDP configured, the rpd process might crash and restart on the new master Routing Engine after a Routing Engine switchover. [PR1155002](#)
- LSPing returns 'routing instance does not exist' when used in vpls routing-instance under logical system. [PR1159588](#)
- If container LSP name and the suffix together are more than 60 characters in length, rpd process might crash during extensive split merge conditions. Its always advisable



to keep them less than 60 characters. The member lsp name is coined in the following manner: <container name>-<suffix name>-<member count> The LSP name can have upto 64 characters. So after putting together the container name, suffix, member-count (could go up to 2 digits), and the 2 hyphens, it should not exceed 64. So container-name and suffix together should not exceed 60 characters. A commit check will be added to throw warning if the name is more than supported character long. [PR1160093](#)

- For BGP-pipe mode OAM the MPLS echo reply will be sent via inet.3 route. [PR1164406](#)
- CE-CE communication over L2VPN(ControlWord Enabled) breaks with chained-composite-next-hop knob. [PR1164584](#)
- Changing maximum-labels configuration under the config hierarchy "protocols mpls interface <>" can cause existing MPLS LSPs to become unusable. When changing this configuration, to make the existing LSPs usable again, the interface in question should be deactivated and reactivated. [PR1166470](#)
- The Output of the command "show mpls container-lsp extensive" was repeating twice. [PR1167533](#)
- In LDP-signaled VPLS environment, other vendor sends an Address Withdraw Message with FEC TLV but without MAC list TLV. The LDP expected that Address Withdraw Message with FEC TLV should always have MAC list TLV. As such, it rejected the message and closed the LDP session. The following message can be seen when this issue occurs: "A@lab> show log messages |match TLV RPD\_LDP\_SESSIONDOWN: LDP session xxx.xxx.xxx.xxx is down, reason: received bad TLV". [PR1168849](#)
- In MVPN scenario, if active primary path goes down, then PLR(Point of Local Repair) needs to send Label Withdraw for old path and new Label Mapping for new path to the new upstream neighbor. In this case, LDP P2MP path may stay in "Inactive" state for indefinite time if an LSR receives a Label Release, immediately followed by a Label Mapping for the same P2MP LSP from the downstream neighbor. [PR1170847](#)
- RPD might crash upon receiving a TLE delete notification arriving during a cleanup sequence. [PR1172567](#)
- When the egress LSR withdraws the label for its egress route, the rlfa nexthop for the ldp route for the egress remains in other routers running rlfs. A routing loop is formed when the rlfa nexthops for some of the router are pointing toward each other. Any traffic for the label route would loop until TTL expires. After the fix, rlfa nexthop with nexthop label alone will not be considered as valid lsp nexthop (primary nexthop). ldp will send label withdraw for the label binding and delete the ldp route to avoid any potential routing loop. [PR1172581](#)
- Multiple RLFA backup gateways (one using spring inner label and other using TLDP label) can get programmed if the given node is PQnode to another node in the network that does not use SPRING RLFA backup for its LDP route, resulting in ECMP among backup next hops. Semantically both gateways provide the same protection path and TLDP based gateway is coming in the way of checking sanity of SPRING backup path. [PR1176489](#)

**Network Management and Monitoring**

- A merge conflict was incorrectly resolved by changing the SNMP trap value of `jnxDomLaneNotifications` to 26. The correct value will always be 25. [PR1145144](#)
- With Junos OS Release 13.3R8/14.1R6/14.1X53-D30/14.2R5/15.1R2/15.1X49-D30 and later, when we configure `fxp0` "master-only" address as source address of SNMP trap, the SNMP trap packets are not sent out after Routing Engine switchover. To restore this issue, we can use "restart snmp" or "delete/set SNMP trap-options". As a workaround, we can use other addresses for snmp trap source. [PR1153722](#)
- Eventd might run out of memory and crash because of excessive kernel logging. [PR1162722](#)
- SNMP statistics extensive command shows incorrect value for "max latency" counter. This PR will correct this behavior. [PR1174029](#)

**Platform and Infrastructure**

- When using the MX2020 platform in a Virtual Chassis (VC) environment, if the Virtual Chassis port (VCP) is located on the local Packet Forwarding Engine whose number is greater than 63 (i.e., VCP is located on local slot 16 or higher), the multicast traffic that should be sent to VCP will be dropped internally due to software issues. As a workaround, please configure the VCP ports on the local chassis (local MX2020) slot 0 to 15, not 16 or higher. [PR1008676](#)
- When the 'enhanced-hash-key services-loadbalancing' feature is used by Trio based line cards, load balancing of flows across multiple service PICs via the source-address across does not work when internal BGP (IBGP) is used to steer traffic to the inside service-interface. For example the operator will see on the stateful firewall that the same source-address has flows across multiple service interfaces. [PR1034770](#)
- Prior to the fix, Juniper VSA length above 2K bytes is not supported. Using authorization parameters above this length would result in wrong authorization setting for the user. [PR1072356](#)
- When one of the "deny-commands" is incorrectly defined in the profile of TACACS+ server, all "deny-commands" regexes will be ignored, which leads to an over-permissive profile without any warning. [PR1078238](#)
- On MX Series-based platforms, when learning the MAC address from the pseudo-IFL (for example, label-switched interface), if the MAC address is aged out in the source FPC where the MAC got learned, due to the delay (around 2 to 3 milliseconds) of MAC address deleting message processed in the source FPC and the egress FPC (destination FPC of the traffic), the MAC address might be deleted first from the egress Packet Forwarding Engine but get added again during these 2-3 millisecond time intervals. (As there is continuous traffic coming on the egress FPC destined to this MAC, the MAC query is generated and sent to the Routing Engine and source FPC. Since the source FPC has not yet processed the MAC-deleted message, it sends the response, so stale MAC will get added on the egress Packet Forwarding Engine.) In this situation, no L2 flooding would occur for the "unknown" unicast (since the MAC address is present on the egress Packet Forwarding Engine). [PR1081881](#)

- With ECMP-FRR enabled, after rebooting the FPC which is hosting some ECMP links, the ECMP-FRR might not work. Clear any BGP sessions (that is the part of ECMP) could help to clear this issue. [PR1101051](#)
- Junos OS defines the SNMP ifXTable (ifJnxInErrors/ifJnxInL3InCompletes) counter as 64-bit width, but it worked as 32-bit width counter. It works as 64-bit width counter after the fix. [PR1105266](#)
- In certain cases, with some events such as disable/enable of links followed by Routing Engine rebooting or GRES enabled switch-over, the following error message could be seen due to a software bug where it doesn't handle an internal flag properly.  
"KERNEL/Packet Forwarding Engine APP=NH OUT OF SYNC: error code 1 REASON: invalid NH add received for an already existing nh ERROR-SPECIFIC INFO:" [PR1107170](#)
- Configuring one group with configuration of routing-instances and applying this group under routing-instances, then the rpd process will crash after executing "deactivating/activating routing-instances" commands. As a workaround, you can avoid using "apply-groups" under routing-instances hierarchy. [PR1109924](#)
- On MX Series platforms, when offlineing the line card (possibly, with any of the line cards listed below), "Major alarm" might be seen due to HSL (link between line card and Packet Forwarding Engine) faults. This fault is non-fatal and would not cause service impact. The line cards that may hit the issue could be seen as below,  
MS-MPC/MS-MIC MIC-3D-8DS3-E3 MIC-3D-8CHDS3-E3-B MIC-3D-4OC3OC12-1OC48  
MIC-3D-8OC3OC12-4OC48 MIC-3D-4CHOC3-2CHOC12 MIC-3D-8CHOC3-4CHOC12  
MIC-3D-1OC192-XFP MIC-3D-1CHOC48 [PR1128592](#)
- On MX Series based line card platform, if FPC offline is performed while FPC is in online progress (online process is at the stage of fabric links training), in very corner scenario, the Routing Engines state is stale and being sent to other existing FPCs, so the traffic forwarding might be affected. [PR1130440](#)
- Doing a file copy from a Routing Engine running Junos OS image to a Routing Engine running Junos OS with Upgraded FreeBSD image fails. [PR1132682](#)
- Too many duplicate ACK messages are generated from Packet Forwarding Engine for TCP control connection with Routing Engine. This could cause: 1. MX-VC DDoS protection violation for VC-control low queue and makds MXVC split. 2. Routing Engine and FPC high CPU utilization. [PR1133293](#)
- With scaled firewall filters attached to interfaces (e.g., 10k+ filters), running the "show configuration" command can cause high CPU of the mgd process. As a workaround, use the "show configuration |display set" command to view the config. [PR1134117](#)
- On XM chip-based line cards (e.g., MPC3/4/5/6, and FPC type 5), in rare situations, when LU or XL chip congestion occurs (e.g., may occur when configuring with more than 4000 entries in the multicast list and large traffic performing replication, please note this is not a realistic configuration), XM chip wedge may occur. [PR1136973](#)
- On MX Series platforms with MX Series base line cards, si interface is configured (i.e., set chassis fpc 1 pic 2 inline-services bandwidth 1g) and service is configured on the si interface. If si ifd is deleted while service is still configured, the FPC might crash. [PR1139348](#)

- When there are additional messages related to FIPS generated during <commit configuration> rpc reply, the xml-tags closing tag <routing engine> may be missed in the reply. [PR1141911](#)
- On ungraceful exit of telnet (quit/shell logout), perm and env files created by pam were not deleted. [PR1142436](#)
- When the CLI command "show pfe statistics exceptions | match reject" is executed CPROD thread in the Packet Forwarding Engine may hog the CPU and result in FPC crash. [PR1142823](#)
- FPC can crash and core due to a missing NULL check. [PR1144381](#)
- When ARP is trying to receive a next-hop message whose size (for example, 73900 bytes) is bigger than its entire socket receive buffer (65536 bytes), the kernel might crash, and the traffic forwarding might be affected. [PR1145920](#)
- In certain affected Junos OS releases, executing the "nhinfo -d" shell command might trigger a kernel panic. This is caused by insufficient buffer space in the routing socket requested by the "nhinfo" utility. [PR1148220](#)
- When the configuration with 6K BFD sessions with 50ms is committed, few BFD sessions may flap while coming up. [PR1148977](#)
- On MX Series platforms with MX Series based line card, inline 6rd with si interface is deployed, if downlink traffic is over ECMP or AE, some traffic might be dropped. [PR1149280](#)
- On MX2000 Series, MPC4 going offline is seen when SFB (Switch Fabric Board) is offlined or removed. This could be caused by the build-up of CDR in ADC which leads to transient packet loss or even getting stuck. The fix prevents line-cards going offline due to transient buildup in ADC. [PR1149677](#)
- When a routing instance is configured with "routing-instances <instance name> routing-options localized-fib", then VPN localization may fail, causing all routes for the affected routing instance to be installed on all Packet Forwarding Engines. [PR1149840](#)
- On MX2010 and MX2020 platform, when error that causes adapter card (ADC)/Switch Fabric Board (SFB) initialization failure occurs (e.g., when Switch Processor Mezzanine Board (SPMB) is bringing up the ADC/SFB which has hardware issue), the SPMB crash occurs. [PR1149910](#)
- When the NTP server address is configured in VRF table and reachable from inet.0 by static configuration (for example, by configuring static/route/next-table/VRF.inet.0), and NTP source-address is configured, the ntpd (the Network Time Protocol daemon running on NTP client) might pick the wrong source-address instead the configured source-address. As a result, NTP server cannot reply the NTP packet back. [PR1150005](#)
- Two interrupts are received from the FPGA on the control board of the MX2010/MX2020 platforms for every i2c transaction triggered from software. Only one is expected. [PR1151674](#)
- FPC may experience blackhole of traffic after lmem data error in private zone. [PR1152026](#)

- During an ISSU upgrade in MXVC environment, linecards may crash causing service impact. When the linecards come up, there may be a next-hop programming issue as a secondary impact and some IFLs may not pass traffic. Affected linecards need to be rebooted to recover from this condition. [PR1152048](#)
- On MX Series routers with Junos OS Release 14.2R5-S1, when we specify a multiservice (ms-) interface to add a timestamp to Real-time Performance Monitor (RPM) probe messages, it will cause the mspmand process crash and the MS-MPC/MS-MIC will keep crashing. As a workaround, configure RPM to perform timestamping either on the Routing Engine (Routing Engine based RPM) or on an installed MPC Packet Forwarding Engine (Inline-RPM). [PR1152785](#)
- Fixed an issue with Inline Jflow where the Observation Domain field in exported IPFIX datagrams was always using the value attributed for LU0 in MPCs with multiple LUs per forwarding-engine. [PR1152854](#)
- With Enhanced LAG mode enabled and sampling configured on AE interfaces, MS-DPC might drop all traffic as "regular discard". Disabling Enhanced LAG mode would avoid this issue. [PR1154394](#)
- The logs CHASSISD\_READBACK\_ERROR are reported on the backup Routing Engine for the non-empty FPCs. [PR1155823](#)
- On MX2000 series platform, when MPC goes down ungracefully, other MPCs in the chassis will experience "destination timeout". In this situation, auto fabric-healing will get triggered due to "destination timeout" condition, which may cause Fabric-Plane reset, even all other MPCs to be restarted in some cases. [PR1156069](#)
- From Junos OS Release 15.1F5 and later, the hidden configuration statement "filter-list-template" will be enabled by default for all firewall filters on MX Series based platforms to use a common program on MX Series-boards for all interfaces that use the same filter list. This can save MX Series board microkernel memory and DMEM memory. The hidden knob "no-filter-list-template" can be configured to disable this behavior. [PR1157079](#)
- Configuring a firewall filter with multiple terms matching either on flexible-match-mask or flexible-match-range might lead to FPC crashing while trying to program the firewall filter and add it to the local table. [PR1157759](#)
- Fixed an issue on where MX Series cards could crash while programming a firewall filter containing flexible-match-mask. [PR1157759](#)
- With Junos OS Release 15.1F2 and later, when inline sampling is enabled on MX Series-based FPC, the srrd (Sampling Route-Record Daemon) process would be created to maintain, collect, and export JFLOW records. On a regular time intervals, the srrd scans through the sampling database for any update/change in the record. In a scaled environment with more route churn, for example 1.14M routes, the scan process might hog CPU for more than 2.5 sec which leads to FPC crash. In some situations, the scan process can run for longer time without causing FPC crash, but it can cause BFD sessions to flap. [PR1158154](#)
- Group names handling process enhancement: one of the core functions was optimized by introducing more efficient pointer comparisons instead of CPU-intensive string ones. [PR1158652](#)

- LU (or XL) and XM chip-based linecard might go to wedge condition after receiving corrupted packets, and this might cause linecard rebooting. [PR1160079](#)
- MPC crashes when "show jnh x hash usage" references incorrect JNH instance (MPC does not crash w/ "show jnh x hash" so it can be used to check). [PR1160697](#)
- The MPC with LU chipset might crash after ISSU. [PR1160748](#)
- NPC cored vpanic in trinity\_firewall\_start\_nh\_get, trinity\_firewall\_add\_and\_check\_internal, trinity\_firewall\_add\_and\_check. This line card core could potentially occur after an ISSU upgrade. [PR1160748](#)
- The following log might be seen when issuing "show jnh x hash usage" from FPC shell: "Feb 10 16:30:37.691 faraday-re0 fpc4 jnh\_hash\_hash\_op\_send(443): PFE 0 Hash TID 1 HASH\_OP not done 0x0300a87601beff03 00000000000000001. Feb 10 16:30:37.692 faraday-re0 fpc4 jnh\_hash\_usage\_get(3217): HASH\_OP failed to count Hash Table Entries for V4 Flows - PFE0." This is due the fact that current timeout of 200msec might not be sufficient for microcode to clean up all flows when issuing this FPC command. Timeout has been increased to 1s. [PR1160775](#)
- The following commit warning may be seen when using configure private and multi-line comments. This causes the commit to not complete: "warning: outgoing comment does not match patch." [PR1161566](#)
- Due to software bug on chassisd, backup CB temperature information is missing on CLI command 'show chassis environment cb' if it's replaced once. [PR1163537](#)
- For MX Series Virtual Chassis with "default-address-selection" configured, when we have a discard route to a specific subnet ( e.g., 10.0.0.0/8 ) with discard next-hop, and at the same time we have more specific routes through other interfaces ( e.g., 10.1.1.1 through xe-0/0/0 ), if a UDP packet is being sent to 10.1.1.1 through xe-0/0/0 while interface xe-0/0/0 flaps or FPC reboots, it might cause kernel crash on both Master Routing Engine in the Virtual Chassis master router (VC-Mm) and Master Routing Engine in Virtual Chassis backup router (VC-Bm). As a workaround, we can disable "default-address-selection" configuration. [PR1163706](#)
- The following log can be seen on MX2020 after one FPC was pulled out and committing the configuration related interface: CHASSISD\_UNSUPPORTED\_FPC: FPC with I2C ID of 0x0 is not supported. [PR1164512](#)
- When two line-cards are taken offline back to back (without delay between issuing the line-card offline command) we are hitting the issue described in the PR (some planes goes to check state). This issue could be prevented by giving a 1-second delay between the offline commands. Workaround would be to offline and online the planes which are in Check state. [PR1164648](#)
- A sonet interface configured as unnumbered BFD session fails to come up. [PR1165720](#)
- Modifying the configuration of a hierarchical policer when in use by more than 4000 subscribers on an FPC can cause the FPC to core and restart. [PR1166123](#)
- There are three issues related to DDOS reported in the PR 1168425. 1) Some policers are configurable, but do not react when disabling them (tunnel-ka aggregate, re-services-v6 capti.v6, syslog aggregate) With the fix all the configurable DDOS

protocol parameter changes will get reflected correctly in Packet Forwarding Engine. 2) Some policers for non-unclassified traffic are non-configurable (control aggregate, mcast-snoop mld, ipsec aggregate, uncls resolve-v4, uncls resolve-v6, uncls filter-v4, uncls filter-v6, tunnel-ka aggregate). These policers are internally deprecated or renamed and not shown on CLI anymore. So any configuration will not come to the Packet Forwarding Engine sides. 3) Some policers are for unclassified traffic are non-zero (mlp unclass, services unclass, radius unclass, ip-frag unclass, gre unclass, re-services unclass, re-services-v6 unclass). We do not have a convention of setting unclassified to 0. Consider this as FAD. [PR1168425](#)

- In Junos 15.1, a customized password prompt that can be sent by a TACACS+ server is not displayed to the user upon login. A usual password prompt "Password: " is displayed instead. The issue is seen when the following conditions are met: 1. Junos OS Release 15.1 without the fix for this PR is used. 2. TACACS+ is used for the user authentication 3. When user logs in, TACACS+ server sends a customized password prompt for this user. For example, this can cause an issue when S/KEY-based one-time password (OTP) authentication is configured for a particular user on the TACACS+ server because the user might be unable to calculate the one-time password as they would not see the key sequence number and the seed provided by the authentication server. [PR1168634](#)
- Because the sequence number in RPM ICMP-PING probes is introduced as 32-bit variable instead of 16-bit, if it increases and reaches the max value 65535, it does not rollover, which might cause all RPM ICMP-PING probes to fail and not succeed any more.. [PR1168874](#)
- In affected release, if user runs the pfe debug command like "show sample-rr eg-table ipv4 entry ifl-index 1224 gateway 113.197.15.66", it will cause the MPC crash. [PR1169370](#)
- Long container elements can have keys which could be very big in size. If the key is more than 256, max key length in Patricia tree, mustd was coring. Now long container, support is added in cdg, long container elements are added in link list, so that they can accommodate any size key. [PR1169516](#)
- Layer 2 protocols might flapp when router was flooded with low priority traffic reaching towards fpc cpu/RE cpu when DDOS protection is disabled. [PR1172409](#)
- On MPC5E/6E/7E/8E/9E/NG linecards, firewall filter of family inet/inet6/vpls configured with non-contiguous prefixes for address matching might fail and cause traffic drop. Using only contiguous prefixes can avoid this issue. [PR1172725](#)
- On all Junos OS platforms, when using RADIUS server, after RADIUS request is successfully sent by Junos device, if the network goes down suddenly, then response sent by the RADIUS server is not received within timeout period. In this scenario, the RADIUS request will be sent again with invalid socket descriptor, which will lead to auditd (provides an intermediary for sending audit records to RADIUS and/or TACACS+ servers) crash. [PR1173018](#)
- On MX2010/2020, MPC/SFB cards do not boot up if single phase AC PSMs are turned ON sequentially with interval even though the PSMs have sufficient remaining power. [PR1176533](#)
- A flow is determined by doing hashing on the packet header. Usually 5-tuple (src/dest IP addresses, IP protocol number, src/dest ports) are used for hashing because a flow



is defined by 5-tuple. This is all fine for TCP/UDP packets. But Layer-3 packets generated by JDSU tester only have Layer-3 header and don't have Layer-4 header. JDSU tester uses the same location as Layer-4 header as packets' sequence number. So MX Series card treats sequence number of JDSU tester packets as Layer-4 header of a packet, hence Junos OS thinks every packet is a single flow and order of different flows is not guaranteed. [PR1177418](#)

- On MX2020/2010, chassisd file rotation on commit check will cause the trace file to be stuck and no other operational chassisd events will be logged until chassisd restart. [PR1177625](#)

### ***Routing Policy and Firewall Filters***

- When a malformed prefix is used to test policy (command "test policy <policy\_name> <prefix>"), and the malformed prefix has a dot symbol in the mask field (e.g., x.x.x.x./24), the rpd process might crash. [PR1144161](#)
- From Junos OS Release 13.2R1, an attempt to commit a configuration with a dangling conditional policy referring a nonexistent/inactive routing-instance will be permitted. If we have a conditional policy referring an active routing-instance, deleting/deactivating this routing-instance and then committing will cause the rpd process crash. As a workaround, we should always make sure that conditional policies are referring active routing-instances. [PR1144766](#)
- interface-routes rib-group import-policy is not in effect to filter prefixes correctly. All direct prefixes could be installed into the secondary route table. [PR1171451](#)

### ***Routing Protocols***

- When configuring router in RR mode (cluster-id or option B MP-eBGP peering), the advertise-external feature will not be applicable in local VRFs due to a different route selection/advertisement process (main bgp.l3vpn.0 vs VRF.inet.0). [PR1023693](#)
- If the command to trace ppm is issued from the FPC shell and a malformed incoming packet (required to be handled by PPM) is in the buffer, the FPC may crash. An example of such a malformed packet would be a multihop BFD packet with an incorrect length (larger than normal). [PR1082878](#)
- This issue is a regression defect introduced in Junos OS Releases 11.4R11, 12.1R10, 12.2R8, 12.3R6, 13.2R4, 13.3R2, 14.1R1. After upgrading to those releases containing the original fix, when there is no export policy configured for the forwarding table to select a specific LSP, whenever routes are resolved over RSVP (for example, due to aggressive auto-bandwidth), the resolver will spend considerable amount of time on the resolver tree, which contributes to the baseline increase in rpd/Routing Engine CPU. [PR1110854](#)
- IGMPv2 working in v2/v1 compatibility mode does not ignore v2 Leave messages received on a bridge-domain's L2 member interface. Moreover, an IGMP snooping membership entry for the respective group at this L2 member interface will be timed out immediately upon IGMPv2 Leave reception, even when there are some other active IGMP hosts attached to this L2 member interface. It might break multicast forwarding for this L2 member interface. [PR1112354](#)
- When two (or more) route target communities of MP-BGP route match to two (or more) route target communities in VRF import policy of a RI, duplicate routing entries



might be installed in the RI. In the output of 'show route table <RI name>.inet.0 detail', two identical routing entries appear with one being marked as 'Inactive reason: Not Best in its group - No difference'. When such duplicate routing information is to be deleted, rpd process process will crash. [PR1113319](#)

- BFD session configured with authentication of algorithm keyed-sha1 and keyed-md5 might be flapping occasionally due to FPC internal clock skew. [PR1113744](#)
- During many types of configuration changes, especially including import policy, BGP has the need to re-evaluate the routes it has learned from peers impacted by the configuration change. This re-evaluation involves re-running import policy to see if there is any changes to the learned routes after applying the new policy. This work is done in the background as part of an "Import Evaluation" job. When BGP is reconfigured a second time, and the "Import Evaluation job" has not completed, it is necessary to re-run the job from the beginning if there's another change to policy or something with similar impact. This state is noted as "Import Evaluation Pending". However, in this case, there was a bug that caused BGP to always enter the pending state upon reconfiguration, regardless of whether relevant changes were made to import or other similarly impactful configuration. The result is that once it is necessary to start re-evaluation of the routes for a peer, even trivial configuration changes that happen too quickly will cause the "Import Evaluation job" to need to run again as a result of the "Pending" flag being set. To avoid the issue, please ensuring that "ImportEval" is not present in a BGP peer's Flags output from the CLI (show bgp neighbor) prior to doing even trivial commits. [PR1120190](#)
- A few seconds of traffic loss is seen on some of the flows when PE-CE interface comes up and PE starts learning 70,000 IPv4 prefixes and 400 IPv6 prefixes from CE during L3VPN convergence. [PR1130154](#)
- In multicast environment, when the RP is FHR (first-hop router) and it has MSDP peers, when the rpf interface on RP changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)
- In multicast environment with Protocol Independent Multicast sparse mode (PIM SM) used, if an upstream router of last-hop router receives the (S,G) SPT join while the shortest-path tree (SPT) is not yet established (only because multicast source is not reachable, a reachable route for SPT which is just not established yet will not cause this issue), when the multicast route gets deleted on the router (e.g., receives the (S,G) prune from downstream PIM router), the router would incorrectly stop forwarding the multicast traffic even if a rendezvous-point tree (RPT) path exists. [PR1130279](#)
- On dual Routing Engine platforms, due to software issue, the OSPF (including both OSPFv2 and OSPFv3) "DoNotAge" bit (e.g., source of LSA has flood-reduction feature enabled) is not mirrored to the backup routing protocol process (rpd). In this situation, after performing nonstop active routing (NSR) switchover, the LSA on new master rpd remains without the "DoNotAge" bit set. Once the LSA reaches OSPF max age, the router will flood LSA purge, hence route flapping might be seen on all routers under the OSPF topology. [PR1131075](#)
- In a situation which BGP is being used in combination with interface's rfp-check; deleted routes may see delay in propagation of BGP withdrawn messages. [PR1135223](#)

- In rare condition, mt tunnel interface flap cause backup Routing Engine core. The exact root cause is not known. While processing updates on the backup Routing Engine (received from master RE), accessing free pointer cause the Core. [PR1135701](#)
- On dual Routing Engine platform with Bidirectional Forwarding Detection (BFD) protocol enabled, after graceful Routing Engine switchover (GRES), the periodic packet management process (ppmd) might crash on the backup Routing Engine due to a software defect. [PR1138582](#)
- When Protocol Independent Multicast (PIM) is used, in very rare condition, if the last hop router (LHR) migrates from (Designated Router) DR to non-DR, repeated routing protocol process (rpd) crash may occur due to patricia tree walk issue.. [PR1140230](#)
- When multicast-only fast reroute (MoFRR) is enabled in PIM or multipoint LDP domain, memory leak will be observed on generation of the multicast FRR next-hops. The leak rate is 8-byte for IPv4 and 12-byte for IPv6 addresses, per FRR next-hop created. Eventually, the rpd process will run out of memory and crash when it cannot honor some request for a memory allocation. [PR1144385](#)
- With NSR configured, when the BFD sessions are replicated on the backup Routing Engine, the master will not send the source address, instead the backup Routing Engine will query the kernel to get the source address. In rare cases, the query might fail, resulting in the source address as all zeros. Later, if a GRES switchover happens, the new master will have this all-zeros source address. When a BFD packet with this source address is send out, the other end will drop the BFD session due to no matching session (source address). [PR1145612](#)
- In the BGP labeled unicast environment, the secondary route is configured with both add-path and advertise-external. If the best route and secondary route are changed in a routing table at the same time, add-path might miss to readvertise the changed route. The old route with the old label is still the last route advertised to one router, instead of updating the advertisement with the new route and new label. So the traffic forwarding might be affected. [PR1147126](#)
- When interface IP MTU is less than 1464 bytes and the number of LSA headers in an OSPF DbD packet is big enough for it to exceed the MTU (i.e., OSPF database contains enough LSAs), unexpected fragmentation of OSPF DbD packets may occur due to incorrect calculation of maximum allowed payload size. [PR1148526](#)
- This core is seen because of incorrect accounting of refcount associated with the memory block which composes the nhid (IRB nh). When the refcount prematurely reaches 0, we released the memory block while it was still referenced from a route. We may see this issue when mcsnoopd becomes a slow consumer of rtsock events generated by rpd (next-hop events in the current case) and messages get delivered in a out-of-order sequence, causing the refcount to be incorrectly decremented. In the testbed where the issue was reported, tracing was enabled for mcsnoopd (for logging all events), causing it to become a slow consumer. However, it may become slow also for other reasons such as processing very high rate of IGMP snooping reports/leaves, which could potentially trigger this issue. [PR1153932](#)
- OpenSSH client software supports an undocumented feature called roaming: if the connection to an SSH server breaks unexpectedly, and if the server supports roaming as well, the client is able to reconnect to the server and resume the suspended SSH

session. This functionality contains two vulnerabilities that can be exploited by a malicious SSH server (or a trusted but compromised server): an information leak (memory disclosure), and a buffer overflow (heap-based). Refer to <https://kb.juniper.net/JSA10734> for more information. [PR1154016](#)

- BGP Monitoring Protocol (BMP) feature is introduced in 13.3R1. When BMP is configured in passive mode and BMP session is closed ungracefully (e.g., No TCP FIN sent), in rare cases, the TCP session might not be cleaned up properly and rpd process crash might be observed during the re-establishment of the previous session. [PR1154017](#)
- In BGP scenario with large scale routing-instances and BGP peers configured, due to a software defect (a long thread issue), BGP slow convergence might be seen. For example, BGP might go down 8-9 seconds after BFD brings down the EBGP session. The rpd slip usually does not hurt anything functionally, but if the slip gets big enough, it could eventually cause tasks to not be done in time. For example, BGP keepalives with lower than 90 seconds hold-time might be impacted. There is no known workaround for this issue, but configuring the knob "protocol bgp precision-timers" can take care of the weak spot like sending BGP keepalives. [PR1157655](#)
- When rib-group copy is done for a route change, the rib-group copy of the secondary route into the destination tables of the copy may not honor maximum-prefixes in some scenarios, such as upon damping changes. The traffic forwarding might be affected. [PR1157842](#)
- Even though no information is actually changed (all ISIS adjacencies remain the same, etc.) when the ISIS LSP is regenerated the different TLVs that compose the LSP might move between the different fragments of the LSP. Although the sum of all the TLVs remains the same the order of the TLVs and their location relative to each fragment might change. The fact that a TVL might move between 2 different fragments might cause issues for ISIS "clients", CSPF for example. [PR1159482](#)
- BFD sessions with keyed authentication might get stuck in init state after system reboot This is only applicable to ACX2100 Platforms [PR1160142](#)
- When a BFD session is configured over an Aggregated Ethernet interface located on a MPC and the MX chassis is set to non-enhanced IP or Ethernet network service mode, with Junos version 15.1F2 or later, the BFD session might be unstable. The workaround is to turn on enhanced-ip mode or disable ppm inline processing. [PR1162716](#)
- Starting from Junos 15.1R1 to 15.1R3, and 15.1F2 to 15.1F4, Junos devices may not be able to establish BGP sessions with legacy router that does not support BGP optional parameters. The reason is that capability of supporting BGP open message fallback to no optional parameter is removed in these releases, which causes "OPEN Message Error (2)" during session setup. [PR1163245](#)
- In BGP scenario with independent domain enabled in a VRF, when configuring a BGP session in a VRF routing instance with a wrong local-as number, some routes might be declared as hidden because of AS path loop. If later configuring the correct AS number as local-as and committing the configuration, those routes might still remain in hidden state. The hidden routes can be released after performing the commands "commit full" or ""clear bgp table <ANY\_VRF>.net.0". [PR1165301](#)

- In L3VPN scenario, feature multipath is configured under [set protocols bgp group] with L3VPN chained CNH under routing-options, the feature multipath does not work for L3VPN routes. [PR1169289](#)
- When clearing IS-IS database, process rpd might crash due to a rare memory de-allocation failure that a task pointer is attempted to be freed twice. In the fix of this issue, the order of referencing the task pointer is being revised to avoid the occurrence of rpd crash. [PR1169903](#)
- PIM bootstrap export policy is not working as expected when there are no PIM neighbors up on the router. [PR1173607](#)
- When we have a route received from different eBGP neighbors, for this specific route, if all BGP selection criteria is matching, we will end up using router ID. As this is eBGP route, so BGP will use active route as the preferred one. Now if this specific route flapped with sequence from the non-preferred to the preferred path, rpd will run the path selection. During rpd path selection we might generate a core file. This issue has no operational impact, also a workaround is available to avoid this issue. [PR1180307](#)
- Next-hop leak could be observed during LSP flap or LSP re-optimization if ISIS is configured in combination with MPLS Traffic Engineering or ISIS te-shortcuts. While this issue does not have an immediate impact, beside higher memory utilization, it could ultimately lead to memory shortage and the inability to program new next-hop structures. [PR1187395](#)

### ***Services Applications***

- In a rare situation in a SIP conversation we might end up in a situation where we have a child conversation whose entry is still present in the parent conversation while the child flow is already deleted. While trying to delete this child flow from the parent conversation, validate if the flow is valid and go ahead with deleting the child flow. [PR1140496](#)
- On MX Series platforms, when using MS-MPC, the "idpd\_err.date" error message is filling var/log. Please refer to KB30743 for details. [PR1151945](#)
- When deleting NAT flow under a race condition, the Service PIC can core. [PR1159028](#)
- In Layer 2 Tunneling Protocol (L2TP) subscriber management environment, the jl2tpd process (L2TP daemon) might crash during cleanup of L2TP tunnel or session after it failed to establish. [PR1162445](#)
- When traffic is flowing through MS-DPC cards Service PIC and there is an active port block and some ports are assigned from that active port block, you should not change the max-blocks-per-address setting to a lower value than the current value since it will cause a core on the cards. Without this fix PR, please don't change the PBA NAT pool configs on the while traffic is going on through Service PIC. You should hold a maintenance window where you can disable the service-set, make the changes and then re-activate the service-set if this change needs to be made. Note even with the fix you still need to hold the maintenance window to deactivate and then activate the service-set for the new configuration change to the max-blocks-per-address setting

to take effect. With the fix we will no longer core when you commit the change.

[PR1169314](#)

- MS-PIC core-dump when MPLS or IPV6 routing updates are received in the PIC.

[PR1170869](#)

### ***Subscriber Management and Services***

- The range for the request-rate statement at the [edit access radius-options] hierarchy level has been extended to 100 through 4000 requests per second. In earlier releases, the range is 500 through 4000 requests per second. The default value is unchanged at 500 requests per second. [PR1033668](#)
- Radius backup accounting queue is used to store Radius records while the Radius server isn't alive. Draining the this queue when the server is reachable again should not log any critical message as this is normal operation. [PR1097491](#)
- When multiple authentication or accounting Radius servers are configured and if one of the servers is down/not-reachable, the Access-Request messages will be queued to the next Radius server no matter if its "max-outstanding-requests" is reached or not. In case that all the Radius servers reached their "max-outstanding-requests", the new requests should be queued to an internal queue, but they are queued to the last Radius server. As a workaround, use only one Radius server or make sure all the Radius servers are reachable. [PR1122703](#)
- When class attribute is changed for a subscriber via CoA, existing subscriber services continue to use the class attribute value at the time when that service was created. Updated class attribute value will take effect for the subscriber and the services created there. When both service and class attributes are present in CoA request, AUTHD first processes the service requests and then processes class attribute. Due to this, accounting starts for requested services do not contain the updated class attribute. [PR1143083](#)
- In normal BRAS environment, if the radius queue is presently full, MX BRAS might stop sending accounting messages and customer might see "Radius result is CLIENT\_REQ\_MAXED\_OUT" in authd log messages. [PR1152052](#)
- shmlogs entries and statistics for AAA daemon (authd) are not visible. [PR1176302](#)

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

- Junoscript traceoptions are available. [PR1062421](#)
- When entering the "restart r" incomplete command in the CLI, the command "restart routing" is executed. It should throw an error like "error: invalid daemon: r". [PR1075746](#)
- When committing a configuration with a very long as-path, in this case the as-path is almost 12000 characters long, the commitd process might crash. The commitd process restart results in a minimal impact on the system. As a workaround, please configure as-path to be less than 4096 characters long. [PR1119529](#)
- While using wildcard with interface like "set groups <group\_name> interfaces <xe> unit <unit>", there is no "disable" option followed. [PR1137377](#)

- When there are two or more sessions accessing the router, and one of the sessions (for example, session 1) is executing commit check in configuration private mode, if another session (for example, session 2) is keeping executing commit and-quit in configuration private mode, because the commit check is not keeping the lock on local Routing Engine for entire session, there is a chance that session 2 will hit a Database opening error. The detailed sequence events are as follows: (1) Session 1: commit check is not keeping the lock on local Routing Engine for entire session, once commit check on local is success, while it asked for lock on other Routing Engine. (2) Session 2: mgd acquired db lock on local Routing Engine. (3) Session 1: once commit check is completed on remote Routing Engine, it does cleanup and deletes the juniper.data+ (created by Session 2). (4) Session 2: juniper.data+ is still in use at local Routing Engine by daemons and daemons start complaining about it and emitting the messages as "Database open failed for file '/var/run/db/juniper.data+' ". [PR1141576](#)
- From Junos OS Release 13.2R1 and later, the commitd process might crash while committing large configurations in a single commit, for example, committing 250k lines of config on top of existing config. This issue is due to a lack of storage space for current and running configurations. [PR1159462](#)

### VPNs

- In a multi-homed source topology in NG-MVPN (applicable to both inter-AS and intra-AS scenario), there are two problems: The first problem is Multicast (S, G) signaling doesn't follow RPF. When the routing table (mvprninstancename.inet0) has two routes, due to the policy configuration, the best route to the source is via the MPLS core, but Multicast (S, G) PIM join and NG-MVPN Type 7 both point to inactive route via local BGP peer. The second problem is when "clear pim join instance NG" is entered, the multicast forwarding entries are wiped out. [PR1099720](#)
- On dual Routing Engine platform with BGP L2VPN and NSR configured, there might be a chance that the block label allocation and deletion for L2VPN is out of order on backup Routing Engine as following: Master rpd follows the below sequeces (which is the correct order): Add Prefix P1 of Label L1 Delete Prefix1 of Label L1 Add Prefix P2 of Label L1 However, on backup rpd, it goes like this: Add Prefix P1 of Label L1 Add Prefix P2 of Label L1 <===== Delete Prefix1 of Label L1 In this situation, backup rpd cannot allocate the label L1 for P2 since L1 is already in use for P1, so it crashes. This occurs in scaling environment (10k L2VPN) where the router has multiple BGP peers and different L2VPN routing-instances are deleted and added back. [PR1104723](#)
- In MVPN scenario, for a race condition, when the forwarding entries go below the threshold, if PIM installs the forwarding entries to reach the forwarding limit, then MVPN will never update the forwarding entry, so it might cause some multicast traffic to be dropped. The correct behavior should be like such: the MVPM should walk the suppress list about entries and try to install the forwarding entries, even if some entries state are moved from 'unsuppress' back to 'suppress'. If there is a PIM installed forwarding entry, then MVPN will be successful in installing the forwarding route. Otherwise, the entry will stay on the suppress list. [PR1144207](#)
- Upon clearing p2mp lsp in dual-home topology, system is adding the same outgoing interface to the (S,G)OIL multiple times and thus duplicate/multiply the amount outgoing traffic. [PR1147947](#)

---

## Resolved Issues: 15.1F5

---

- [General Routing on page 127](#)
- [Class of Service \(CoS\) on page 132](#)
- [Forwarding and Sampling on page 132](#)
- [High Availability \(HA\) and Resiliency on page 132](#)
- [Infrastructure on page 133](#)
- [Interfaces and Chassis on page 133](#)
- [Layer 2 Features on page 136](#)
- [MPLS on page 136](#)
- [Network Management and Monitoring on page 137](#)
- [Platform and Infrastructure on page 137](#)
- [Routing Policy and Firewall Filters on page 139](#)
- [Routing Protocols on page 140](#)
- [User Interface and Configuration on page 141](#)
- [VPNs on page 141](#)

### **General Routing**

- With "chassis maximum-ecmp 64" configured, when there is a route having 64 ECMP LSP next-hops and CoS-based forwarding (CBF) is enabled with 8 forwarding class (64\*8=512 next-hops), not all next-hops will be installed on Packet Forwarding Engine due to crossing the boundary in the kernel when number of ECMP next-hops is large than 309. [PR917732](#)
- On an MX Series Virtual Chassis platform, when we restart one or both of the standby Routing Engines, the log message "ksyncd\_select\_control\_plane\_proto: rhost\_sysctlbyname\_get: No such file or directory" might be observed as the ksyncd daemon attempts to select a communication protocol (UDP/TCP). After several tries, it will fall back to TCP and proceed as normal. [PR945925](#)
- On MX240/480/960/2010/2020 platforms with Junos OS Release 15.1R1 or later, the process health monitor process (pmond) is not available on the Routing Engine. The mspmond process on the MS-MIC/MS-MPC tries to connect to the pmond process on Routing Engine continuously but fails. This will result in additional traffic between the MS-MIC/MS-MPC and the Routing Engine, causing high CPU utilization. [PR1014584](#)
- On all MX Series routing platforms with BGP configured to carry flow-specification route, in case of deleting a filter term and policer, then add the same term and policer back (it usually happens in race condition when adding/deleting/adding the flow routes), since confirmation from dfwd for the deleting policer might not be received before attempting to add the same policer, the rpd would skip sending an add operation for it to dfwd. As a result, when the filter term is sent to dfwd and tell it to attach to the policer, dfwd had already deleted the policer, and since rpd skipped re-adding it, dfwd will reject the attach filter with policer not found error and rpd will crash correspondingly. [PR1052887](#)



- When flag is specified under ipsec-vpn traceoptions to trace IPsec operations, no message is logged to the specified trace file as expected. The issue impacts on debug capability only. [PR1073705](#)
- When configuring the large-scale firewall filter (e.g., with 10K terms on input/output) on either FPC5 or MPC3/4/5/6, traffic drop might occur due to allocation limits. [PR1093275](#)
- When route convergence occurred, the new gateway address is not updated correctly in inline-jflow route-record table (route-record table is used by sampling), and the sampling traffic forwarding might be affected, but normal routing would be not affected. [PR1097408](#)
- Fragmenting a special host outbound IP packet with an invalid IP header length (IP header length is greater than actual memory buffer packet header length) can trigger NULL mbuf accessing and dereferencing, which might lead to a kernel panic. [PR1102044](#)
- During NSR Routing Engine switchover, there might be a control plane black window for inline BFD causing the BFD session to flap. This is a day-1 issue, and tuning the retrans timer would solve the problem. But since these timers have to meet RFC compliance, we cannot really do that. Today we have the retrans timer as 1000 milliseconds. The workaround would be to configure a higher retrans timer value. [PR1105980](#)
- On MX240/480/960 Series routers with MS-DPC, customer is running BGP over IPsec. This BGP session has a BFD session tied to it. The BGP session is up but the BFD session remains in INIT state. The issue might be seen with any service configured with multi-hop BFD enabled. Traffic forwarding will not be affected. [PR1109660](#)
- This issue is a regression defect introduced in Junos OS Release 11.4R11, 12.1R10, 12.2R8, 12.3R6, 13.2R4, 13.3R2, 14.1R1. After upgrading to those releases containing the original fix, when there is no export policy configured for the forwarding table to select a specific LSP, whenever routes are resolved over RSVP (for example, due to aggressive auto-bandwidth), the resolver will spend considerable amount of time on the resolver tree, which contributes to the baseline increase in rpd/Routing Engine CPU. [PR1110854](#)
- On MX Series routers containing multiple Packet Forwarding Engines such as MX240/MX480/MX960/MX2010/MX2020, with MPC3E/MPC4E/MPC5E/MPC6E cards, if the routers have GRE decap, then certain packet sizes coming via these line cards, at very high rate can cause these line cards to exhibit a lockup, and one or more of their Packet Forwarding Engines corrupt traffic towards the router fabric. [PR1117665](#)
- During the LSP switchover, the hiwatermark might get set to an unexpectedly high value. The issue happens due to an incorrect reference point taken while calculating the Max avg BW in the last interval, and this results in an incorrect Highest Watermark BW in the autobandwidth stats. [PR1118573](#)
- MX Series router acting as an L2TP access concentrator (LAC) might not recognize the MLPPP protocol field (0x003d) in the inbound PPP packet from the customer premise equipment (CPE) and could disconnect the session not respecting idle-timeout. The traffic forwarding might be affected. [PR1123233](#)



- On MX Series platforms, the MS-MPC crash might occur. The exact trigger of the issue is unknown; normally, this issue might happen over long hours (e.g., within a week) of traffic run (e.g., running HTTP/HTTPS/DNS/RTSP/TFP/FTP traffic profile). [PR1124466](#)
- In an EVPN scenario, the EVPN route table between the master Routing Engine and backup Routing Engine would be different (unused garbage routes will appear) once Routing Engine switchover (e.g., by rebooting the "old" master Routing Engine or performing a graceful Routing Engine switchover) is performed, which might cause a kernel crash on the new master Routing Engine in some cases. [PR1126195](#)
- When Junos OS devices use the Link Layer Discovery Protocol (LLDP), the command "show lldp neighbors" displays the contents of PortID type, length, and value (TLV) received from the peer in the field 'Port Info', and it could be the neighbor's port identifier or port description. A Junos OS CLI configuration statement can select which "interface-name" or "SNMP ifIndex" to generate for the PortID TLV, so we do not have any problem as long as two Junos OS devices are connected for LLDP, but we might have an interoperability issue if another vendor device that can map the configured 'port description' in the PortID TLV is used. In this case, Junos OS displays the neighbor's PortDescription TLV in the Port info field, and if the peer sets the port description whose TLV length is longer than 33 bytes (included), Junos OS is not able to accept the LLDP packets and discards the packets as errors. The PortID TLV is given as : "the port id tlv length = port description field length + port id subtype(1B)". [PR1126680](#)
- In multi-homing Ethernet VPN (EVPN), if there are two loopback addresses and the router-id and the primary loopback addresses are different on the designated forwarder (DF) PE, when the link between CE and DF PE is down, the Type 4 route of old DF not deleted properly from the backup PE, causing the new DF election failure. The traffic forwarding will be affected. As a workaround, we should configure a single primary loopback address and remove "router-id" configuration statements on both multi-homing PEs. [PR1126875](#)
- On MX Series routers with MS-MIC (or possibly, MS-MPC is affected as well), changing the configuration of sampling input parameters, such as "rate" under forwarding-options, is not reflected without restarting the line card. [PR1131227](#)
- CLI output of "clear services sessions" gives an impression to the user that the session is marked for deletion in case of delayed delete, but the XML output "clear services sessions|display xml" of the above command says "session removed." Ideally both should convey the same message to the user. The changes have been made to make sure CLI and XML information given to the user in sync. [PR1132006](#)
- When customers do changes under "protocol router-advertisement interface X" (such as changing timers, etc.), they expect that a commit would trigger a new router-advertisement being sent out to notify hosts about configuration changes. However, this does not seem to be the case, unfortunately. It makes the router information to expire on hosts and causes obvious loss of connectivity for the hosts. [PR1132345](#)
- The subscribers login rate could be degraded when IGMP/MLD is enabled on the dynamic demux interface. [PR1134558](#)
- MXVC-Same subnet VC-heartbeat polling failed to recover. [PR1136119](#)

- On MX Series platforms with MIC3-3D-1X100GE-CFP, after unified in-service software upgrade (ISSU), the Junos OS upgrade is successful, but the 100GE port will be down, and the traffic forwarding will be affected. [PR1136269](#)
- While checking JNH pool usage on MPC cards, the error listed below might be logged due to the fact that those cards do not have physical bulk DMEM. This has been addressed by adding an extra check in the code before fetching the data from the card.  

```

NPC4(faraday-re1 vty)# sh jnh 0 pool usage EDMEM overall usage:
[NH////////|FW////////|CNTR////////|HASH/////|ENCAPS////|-----]
-----] 0 4.0 8.0 14.0 18.0 22.0 32.0M Next Hop
[*****|-----|RRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRR] 4.0M (36% |
64%) Firewall [|-----|RRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRRR] 4.0M
(1% | 99%) Counters
[*****|-----]
6.0M (35% | 65%) HASH
[*****] 4.0M (100% | 0%)
ENCAPS [*****] 4.1M (100%
| 0%) Shared Memory - NH/FW/CNTR/HASH/ENCAPS
[-----] 10.0M
(0% | 100%) NPC4(faraday-re1 vty)# [Feb 4 09:39:55.377 LOG: Err]
jnh_partitions_show_usage_helper(8835): Error from (PFE 0)
jnh_partitions_get_usage_stats. PR1136481

```
- In a IGMP oversubscriber environment with the configuration statement "remove-when-no-subscribers" configured, after performing graceful Routing Engine switchover, subscribers with multicast joins cannot re-login when the subscriber logs out before it sends an IGMP leave in the new master. [PR1136646](#)
- Insufficient time to allow an MPC5/MPC6 card to lock on the clocking source during FPC boot time might cause the Major Alarm raised due to "PLL Error." [PR1137577](#)
- On MS-MIC, TCP session Up/Down causes JSERVICES\_NAT\_\* and JSERVICES\_SESSION\_\* messages even though severity level "none" is configured for services. [PR1137596](#)
- When successive back-to-back commits are performed on a scaled configuration, there could be a timeout or a delay in completing the commit check operation. [PR1139206](#)
- Using several Junos OS 15.1 daily builds post 15.1Rx, the IPFIX flow data for ICMPv6 packets, value of ICMPv6 type, and code (icmpTypeCodeIPv6) are wrongly stored as L4 source port (sourceTransportPort). This issue is observed on both MPC7E and MPC3E. This issue now fixed and committed to 15.1F5. [PR1139986](#)
- JNH periodically attempts to recover memory no longer in use. Recently, when firewall address space was expanded to 16M, a side effect was triggered -- memory recovery was extended to 16M as well. On the Hercules line card, Firewall does not use a small block of IDMEM, causing JNH to attempt the return of the unused memory. There is no mechanism for recovery of IDMEM, therefore, this message is displayed. Excepting the syslog impact, there is no further effect on the line card. [PR1140021](#)
- In Junos OS Release 15.1F4, "show chassis environment" "Routing Engine 0 CPU" does not show Routing Engine CPU temperature as Junos OS Release 15.1F4 does not have

the fix to make available the CPU temperature in Junos OS from HOST. "Routing Engine 0 CPU" instead shows the maximum of the Routing Engine inlet and exhaust sensors reading. [PR1140187](#)

- After removing a child link from AE bundle, the AE interface statistics in the SNMP MIB might show a spike. [PR1140533](#)
- From Junos OS Release 14.1R4, 14.2R3, 15.1, and later, when a firewall filter is applied to NG-MPC, after system reboot, the Routing Engine might go into amnesiac mode. [PR1141101](#)
- The unified in-service software upgrade (ISSU) never works fine when hyper-mode feature is enabled on enhanced MPCs such as MPC3E, MPC4E, MPC5E, and MPC6E. Prior to Junos OS Release 15.1R3/15.1F4/14.1X51-D60, both ukernel image and ucode image are getting upgraded to normal mode; while from those releases and later, traffic will be dropped on Enhanced MPCs, the issue can be recovered by rebooting enhanced MPCs. [PR1144648](#)
- In certain affected Junos OS releases, executing "show arp" or "clear arp" might trigger a kernel panic. This is caused by insufficient buffer space in the routing socket requested by the "arp" utility. [PR1145920](#)
- When a route in VRF has a indirect next hop, and the indirect next hop is pointing to an interface which is using an un-numbered address, then the route in the VRF table might be stuck in the KRT queue. [PR1147776](#)
- On MX Series platforms, in a multicast subscriber management environment (e.g., IGMP is configured for subscribers in dynamic profile), when nonstop active routing (NSR) is enabled, if the routing protocol process (rpd) is busy or hundreds of multicast groups are active (e.g., 250), the missing multicast entries issue might be seen after performing Routing Engine switchover twice or more (i.e., the first Routing Engine switchover works fine, and the issue might occur from the second switchover and onward). As a workaround, this issue could be avoided by issuing the CLI command "restart smg-service" on the backup Routing Engine after every switchover. [PR1149065](#)
- When a routing instance is configured with "routing-instances <instance name> routing-options localized-fib," then VPN localization might fail, causing all routes for the affected routing instance to be installed on all Packet Forwarding Engines. [PR1149840: This issue has been resolved.](#)
- Commit error after attempting to delete all guaranteed rates on all traffic-control-profiles associated with demux0 [edit] lab@mx480-J12\_09# commit re0: [edit class-of-service interfaces] 'demux0' IFL excess rate not allowed on interface (demux0), please specify guaranteed rate on at least one IFL error: configuration check-out failed. [PR1150156](#)
- Routers using inline Layer 2 services might experience fabric degradation and FPC restart. This problem is amplified by fragmented and out-of-order packets. This log entry might be seen during the error state: Host Loopback:HOST LOOPBACK WEDGE DETECTED IN PATH ID 0. [PR1153750](#)

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

- On MX104 platforms, when applying the "rate-limit" and the "buffer-size" on the logical tunnel (lt-) interface on the missing MIC (not inserted on MPC), a commit failure with error message would occur. As a workaround, this issue could be avoided by applying the "rate-limit and "buffer-size" on the inserted MIC, then commit. [PR1142182](#)

***Forwarding and Sampling***

- On MX80 and MX104 platforms, applying a firewall filter with an MX Series specific match condition will raise the following warning message: Filter <filter\_name> is MX Series specific; will not get installed on DPCs for interface <interface\_name>. This warning message is needed for the other modular-type MX Series platforms since they can have DPC and MPC mixed. But the message is not needed for MX80 and MX104 platforms since they only have the MX series-based Packet Forwarding Engine. Although the warning message indicates that the relevant firewall filter is not installed, the firewall filter is correctly installed into the Packet Forwarding Engine. Thus, the user can ignore the message in case it is logged on MX80 and MX104 platforms. [PR1138220](#)
- On MX Series-based platforms, in race condition, when using the policer that has the configuration statement "bandwidth-percent" configured (e.g., set firewall policer XXX if-exceeding bandwidth-percent 80), if the logical interface (IFL) bandwidth change and the filter bind message arrive at the Packet Forwarding Engine out of order (e.g., when changing the bandwidth of the IFL or rebooting the FPC), the "bandwidth-percent" policer might end up using physical interface (IFD) bandwidth for "bandwidth-percent" computation. [PR1154034](#)

***High Availability (HA) and Resiliency***

- On MX240/480/960/2010/2020 platforms with Junos OS Release 15.1R1 and later, in a high-scale scenario (e.g., there are 4 million routes or more), the connection between Routing Engine and the FPC(s) might flap after performing graceful Routing Engine switchover (GRES). The other symptoms are intermittent packet drops between the Routing Engine and FPC during regular operation without performing GRES and scaled scenario. [PR1146548](#)
- Unified ISSU between Junos OS Release 15.1F3 and earlier to Junos OS Release 15.1F4, and ISSU between releases from Junos OS Release 15.1F4 to Junos OS Release 15.1F5 will result in a core dump and could lead to [PR1161491](#). The same might happen when the ISSU is done from 14.2R4/R5 to 15.1F5 only. This issue happens due to an inconsistency in port numbering between two port types in the releases. There could be other consequences due to this issue in the upgraded release that might hamper functionality on some types of ports only. [PR1161491](#)
- MXVC: Unified ISSU failed after all FPC upgraded, TCP connection to kernel was dropped due to invalid IPC type 20. [PR1163807](#)
- When configuring the "nonstop-routing" under one group and apply this group to routing-options configuration hierarchy, sometimes the NSR does not work. As a workaround, please configure the "nonstop-routing" directly under the routing instance hierarchy. [PR1168818](#)

### **Infrastructure**

- In scaling setup (in this case, there are 1000 VLANs, 1000 bridge domains, 120 IRB interfaces, 120 VRRP instances, BGP, and IGP), if the routing protocols are deactivated and activated, there might be a chance that the pending route stats are not cleaned up, which will cause the stats infrastructure to have stale pointers and lead to memory corruption in socket layers. The system might go to the db prompt because of this. All the traffic going through the router will be dropped. [PR1146720](#)
- In Legacy, if the Routing Engine switchover on rpd crash configuration statement is enabled, the switchover is instigated in the Routing Engine kernel BEFORE the rpd core is created. In Occam, it is done AFTER. This creates an issue in scaled setups where the size of the rpd core, and therefore the time to create it, takes a lot longer. An Occam FreeBSD change has been made as part of this PR patch to adopt the same behavior as Legacy. [PR1164139](#)
- Support for fast Routing Engine mastership switchover when master Routing Engine crashes was broken in Occam releases up to and including 15.1R3. The patch for this PR addresses support for this feature in Occam in 15.1R4. [PR1167385](#)
- Issue is specific to Occam-based images and is a very hard to induce. The issue occurs when virtual memory is low, and the FreeBSD OS tries to free up memory by invoking the vm\_lowmem event. In a multi-core environment, multiple kernel threads could encounter the low virtual memory condition at the same time, and so the vm\_lowmem event could be invoked within the context of multiple kernel threads concurrently. Some of the protocols in the Junos networking stack register handlers against this event and two of these, clnp\_drain() & tcp\_drain(), were not SMP safe, which caused data corruption. clnp\_drain() & tcp\_drain() have now been made SMP safe; all other such handlers in the Junos networking stack were already SMP safe. [PR1182958](#)

### **Interfaces and Chassis**

- On dual Routing Engine platforms, when adding the logical interfaces (IFLs) and committing, due to the device control process (dcd) on the backup Routing Engine might fail to process the configuration and keep it in the memory. In some cases (not happening all the time), it might be observed that the memory of the dcd keeps increasing on the backup Routing Engine. [PR1014098](#)
- jnxBoxDescr is reworded for MXVC to replace the platform type with a more general representation that replaces the specific member platform type with "Virtual Chassis." Old virtual chassis text example: jnxBoxDescr.0 = member0 Juniper MX240 Internet Backbone Router. New virtual chassis text example: jnxBoxDescr.0 = member0 Juniper MX Virtual Chassis Internet Backbone Router. NOTE: The MIB design for jnxBoxAnatomy "top-level" chassis information works properly for a standalone chassis, but does not fully represent virtual chassis multi-member configurations because it is capable of providing information for only one physical chassis. (The remainder of the jnxBoxAnatomy MIB "containers" properly support the inventory of a multi-member configuration.) MX virtual chassis provides another MIB, jnxVirtualChassisMemberTable, to supply the equivalent "top-level" information. [PR1024660](#)
- MS-DPC might crash when allocating chain-composite next hop in an enhanced LAG scenario. [PR1058699](#)

- During failure notification state machine, CFM does not correctly transit from DEFECT CLEARING state to RESET once the error indication has been cleared. As a consequence, all the forthcoming errors will be considered post errors and will be reported right away without incurring the fngAlarmTime. This is a cosmetic problem. [PR1096346](#)
- When DHCP subscribers are terminated at specific routing-instances and the interface stack is IP demux over vlan-subinterface over AE interface, there might be a memory leak in the kernel AE iffamily when subscribers log in/log out. [PR1097824](#)
- The following CLI configuration statement needs to be used for the CFM session to work: "set chassis aggregated-devices disable-lag-enhanced." Enhanced-lag is enabled by default in the system when the system is configured with enhanced-ip. CFM is not supported with enhanced-lag at present. [PR116826](#)
- If two redundant logical tunnel (rlt) sub-interfaces are configured in the same subnet and in the same routing-instance, a sub-interface will be down (this is expected), but if the sub-interface is removed from the routing-instance later, after disabling and enabling the rlt interface, a sub-interface might remain in the down state unless you remove the configuration of the rlt interface and then do a rollback. [PR1127200](#)
- In the dual Routing Engines scenario with fast-synchronize configuration, an interface is added as part of an interface-set configuration. When the interface is deactivated, as fast-synchronize is configured, the commit check operation is not executed on the backup Routing Engine. Due to this, the commit check error is not caught and the commit operation is forwarded to the backup Routing Engine, also resulting in error conditions at run time. [PR1128038](#)
- MXVC-specific behavior for SNMP walk of jnxOperating\* containers was divergent from the physical MX Series. Returned to vergence. [PR1136414](#)
- %DAEMON-3-CHASSISD\_I2C\_WRITE\_ERROR: i2cs\_write\_reg: write error for group 8 at address 0x49, offset 32 %DAEMON-3-CHASSISD\_I2CS\_READBACK\_ERROR: Readback error from I2C slave for FPC 1 ([0x11, 0x42] -> 0x0) - The above errors represent transient communication issues between system components. In certain cases, these can be service impacting. Enhancements have been made for better handling of such error conditions. [PR1139920](#)
- When micro Bidirectional Forwarding Detection (BFD) sessions are configured for link aggregation group (LAG), the device control process (DCD) acts as the client to the micro BFD session. In order to monitor the connection between client (DCD) and server(BFD), the client needs to exchange keepalive hello packets with the server. To send hello packets, DCD needs to move out of IDLE phase to CONFIG\_BFD phase, which is the reason for the following log messages: "dcd.c:585 dcd\_new\_phase\_if\_idle() INFO : Current phase is IDLE, going to phase CONFIG\_BFD usage.c:75 dcd\_trace\_times() INFO : Phase Usage for IDLE : user 0.001 s, sys 0.000 s, wall 60.019 s dcd.c:717 dcd\_new\_phase() INFO : New phase is CONFIG\_BFD usage.c:75 dcd\_trace\_times() INFO : Phase Usage for CONFIG\_BFD : user 0.000 s, sys 0.000 s, wall 0.000 s dcd.c:717 dcd\_new\_phase() INFO : New phase is IDLE". There is no functionality impact; however, these messages might flood the logs. As a workaround, we can filter out these messages from being written to the log file according to this KB article: <https://kb.juniper.net/InfoCenter/KB9382>. [PR1144093](#)

- In MX-VC or VRR platforms running Junos OS Release 15.1 built before about February 2016, the following cosmetic warning message will be displayed upon commit: "[edit] 'chassis' warning: WARNING: MPC reboot or chassis reboot is required to use MIC aware dynamic power management feature on already plugged-in MPCs". [PR1144295](#)
- During a VRRP configuration change involving IP address change and/or VRRP configuration change while retaining same group ID, a race condition might occur causing vrrpd crash. [PR1145170](#)
- When using MX Series platforms as Layer 2 Tunnel Protocol (L2TP) L2TP access concentrator (LAC), if login/logout tunneled PPPoE subscribers over an extended period (e.g., login/logout 16K subscribers for 24- 48 hours), kernel crash may occur due to next-hop issue. [PR1150316](#)
- In affected releases, the following cosmetic alarms are seen after reseating the clocking cables: 2015-11-13 05:22:56 UTC Major CB 0 External-A LOS 2015-11-13 05:22:56 UTC Major CB 0 External-B LOS. [PR1152035](#)
- Remove MX Series from sending LCD halt message. [PR1153219](#)
- The outbound PPPoE control packets sourced from PPPoE daemon, such as PADO, error PADS and PADT, are assigned to queue 0 instead of queue 3. [PR1154070](#)
- SONET interface on MIC-3D-IOC192-XFP does not count input error correctly. While hardware counts framing error, runts and giants but input error in 'show interface extensive' command reports runts and giants only. [PR1154268](#)
- Customer might see errors when doing 'show interface interface-set queue <if set>' for a pure numeric interface-set name. router> show interfaces interface-set queue 803 error: cannot decode interface name `803': invalid device name. [PR1154667](#)
- When the master Routing Engine in the Virtual Chassis master router (VC-Mm) runs with high CPU (e.g., 99 % CPU utilization), after a global/local switchover, the new master Routing Engine might relinquish its mastership during high CPU conditions. But the Virtual Chassis protocol role is not changed properly after the kernel relinquishes the mastership, causing dual master Routing Engines on this member router. [PR1156337](#)
- MX Series Routing Engine high CPU due to stale ifmon process. [PR1162521](#)
- Internal timing for bringing FPCs online is extended for MX 2020/2010 systems to accommodate longer initialization times for fabric and FPCs. [PR1164147](#)
- CLI commit warning is replaced by syslog warning message when limited-ifl-scaling configuration statement is configured. Warning message text remains the same. [PR1165357](#)
- If an interface configured with VRRP is removed from a routing-instance to global, or from global to a routing-instance, the IFLs of that interface will be deleted and recreated. In ideal cases as the interface gets deleted, VRRP should move to bringup state; when the interface is created again, VRRP goes to previous state. After this, VRRP should get VIP addition notification from kernel and update VRRP state and group id for VIP. However, in race conditions, VRRP might get VIP addition notification from kernel even before the interface creation event happens. If so, VRRP will never be able to update proper VRRP state and group id. So the VIP will reply for the ARP with an incorrect



MAC ending with "00", while the correct MAC should end with the groups id configured. [PR1169808](#)

- When upgrading Junos OS software on RE1, if at the time, RE1 is the "master Routing Engine", both Routing Engines will be in "backup" state, resulting in losing remote connectivity, and all interfaces. Only "console" access will be available at this time. [PR1172729](#)
- Commit check may exit without providing correct error message and causing dcd exit. The only known scenario to trigger this issue is to configure a IPv6 host address with any other address on the same family. [PR1180426](#)

### **Layer 2 Features**

- In BGP-based VPLS scenarios, changing the configuration of a VPLS mesh group might cause rpd core. FPC reboot might also be seen during the rpd core. [PR1123155](#)
- In a VPLS scenario, when "\$junos-underlying-interface-unit" is configured in "dynamic-profiles" hierarchy, which is then implemented in a routing-instance, upgrade/commit will fail with the following error message: Parse of the dynamic profile (<dynamic\_profile\_name>) for the interface: \$junos-interface-ifd-name and unit: \$junos-underlying-interface-unit failed. [PR1147990](#)
- For router equipped with following line cards: T4000-FPC5-3D MX-MPC3E-3D MPC5E-40G10G MPC5EQ-40G10G MPC6E MX2K-MPC6E, if the router is working as VPLS PE, due to MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)
- From Junos OS Release 13.2R1 and later, the rpd process might crash when adding/deleting Virtual private LAN service (VPLS) neighbors in a single commit. For example, a primary neighbor is changed to become the backup neighbor. [PR1151497](#)
- There is a bug in code of handling the redistribution of PPM (periodic packet management) Transmit and Adjacency entries for LACP, when the Interface entry is in pending distribution state. This issue might cause a ppmd crash after graceful Routing Engine switchover. [PR1116741](#)

### **MPLS**

- In MPLS scenarios, removing the "family mpls" configuration from an outgoing interface might cause inet and/or inet6 next hops associated with that interface to unexpectedly transit to dead state. Even adding back "family mpls" cannot restore it. [PR1067915](#)
- For advertising IPV6 packets over the MPLS GRE tunnel, the IPV6 address gets stuck in KRT queue. [PR1113967](#)
- When a PLR is a non-Juniper router, Juniper ingress node might stay on the bypass tunnel and ignore the CSPF result. [PR1138252](#)
- When a link fails on an RSVP LSP which has link-protection or node-link-protection configured, the PLR (point of local repair) will initiate a bypass LSP and the RSVP LSP will be tunneled on this bypass LSP. However, if now the bypass LSP is brought down because there is a link failure on it, the PLR might only send out a session\_preempted PathErr message to the upstream node without sending a ResvTear message. Hence the ingress node does not receive a ResvTear message and the RSVP LSP is not



immediately torn down. The RSVP LSP will remain UP for more than 2 minutes until the RSB (Resv state block) on the ingress's downstream node gets timed out and it sends a ResvTear message to the ingress. [PR1140177](#)

- There is no entropy label for LDP route in a scenario of LDP tunneling across a single hop RSVP LSP with label 0 (explicit-null) used. As workaround, either remove LDP tunneling or RSVP explicit-null will resolve the issue. [PR1142357](#)
- This issue is related to inter-op between multi-vendor scenario. This fix will add sub-object RRO which will help change of label during FRR active scenario. [PR1145627](#)
- With NSR enabled and LDP configured, the rpd process might crash and restart on the new master Routing Engine after a Routing Engine switchover. [PR1155002](#)

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

- A merge conflict was incorrectly resolved by changing the SNMP trap value of jnxDomLaneNotifications to 26. The correct value will always be 25. [PR1145144](#)

### ***Platform and Infrastructure***

- When using MX2020 platforms in a Virtual Chassis (VC) environment, if the Virtual Chassis port (VCP) is located on the local Packet Forwarding Engine whose number is greater than 63 (i.e., VCP is located on local slot 16 or higher), the multicast traffic that should be sent to VCP will be dropped internally due to a software issue. As a workaround, please configure the VCP ports on the local chassis (local MX2020) slot 0 to 15, not 16 or higher. [PR1008676](#)
- On MX Series-based platforms, when learning the MAC address from the pseudo-IFL (for example, label-switched interface), if the MAC address is aged out in the source FPC where the MAC got learned, due to the delay (around 2 to 3 milliseconds) of MAC address deleting message processed in the source FPC and the egress FPC (destination FPC of the traffic), the MAC address might be deleted first from the egress Packet Forwarding Engine but get added again during these 2-3 milliseconds time intervals (As there is continuous traffic coming on the egress FPC destined to this MAC, the MAC query is generated and sent to the Routing Engine and source FPC. Since the source FPC has not yet processed the MAC-deleted message, it sends the response, so stale MAC will get added on the egress Packet Forwarding Engine). In this situation, no L2 flooding would occur for the "unknown" unicast (since the MAC address is present on the egress Packet Forwarding Engine). [PR1081881](#)
- With ECMP-FRR enabled, after rebooting the FPC which hoisting some ECMP links, the ECMP-FRR might not work. Clearing any of BGP sessions (that is the part of ECMP) could help to clear this issue. [PR1101051](#)
- For IPv6 packet with "no next header" in Hop-By-Hop header, if the Hop-By-Hop header length field value is large than 112, the router will drop such packet and log the following error: "PPE PPE HW Fault Trap: Count 105, PC 60ce, 0x60ce: ipv6\_input\_finished\_parsing LUCHIP(3) PPE\_10 Errors lmem addr error". [PR1130735](#)
- NTP.org published a security advisory for thirteen vulnerabilities in NTP software on Oct 21st, 2015. These vulnerabilities may allow remote unauthenticated attackers to cause Denial(s) of Service(s), disruption of service(s) by modification of time stamps being issued by the NTP server from malicious NTP crafted packets, including

maliciously crafted NTP authentication packets and disclosure of information. This can impact DNS services, as well as certificate chains, such as those used in SSL/https communications and allow attackers to maliciously inject invalid certificates as valid which clients would accept as valid. [PR1132181](#)

- Too many duplicate ACK messages are generated from Packet Forwarding Engine for TCP control connection with Routing Engine. This could cause: 1. MX-VC DDoS protection violation for VC-control low queue and makds MXVC split. 2. Routing Engine and FPC high CPU utilization. [PR1133293](#)
- With scaled firewall filters attached to interfaces (e.g., 10k+ filters), running "show configuration" command can cause high CPU of the mgd process. As a workaround, we can use the "show configuration | display set" command to view the configuration. [PR1134117](#)
- On ungraceful exit of telnet (quit/shell logout), perm and env files created by pam were not deleted. [PR1142436](#)
- When the CLI command "show pfe statistics exceptions | match reject" is executed, CPROD thread in the Packet Forwarding Engine may hog the CPU and result in FPC crash. [PR1142823](#)
- Sometimes Inline jflow incorrectly reports SNMP index of internally generated LSI interface instead of SNMP Index of Actual outgoing interface in Information Element ID 14 in VPLS IPFIX flow records. [PR1143699](#)
- In certain affected Junos OS releases, executing "nhinfo -d" shell command might trigger a kernel panic. This is caused by insufficient buffer space in the routing socket requested by the "nhinfo" utility. [PR1148220](#)
- When the configuration with 6K BFD sessions with 50ms is committed, few BFD sessions may flap while coming up. [PR1148977](#)
- On MX2010 and MX2020 platform, when error that causes adapter card (ADC)/Switch Fabric Board (SFB) initialization failure occurs (e.g., when Switch Processor Mezzanine Board (SPMB) is bringing up the ADC/SFB which has hardware issue), the SPMB crash occurs. [PR114991](#)
- When the NTP server address is configured in VRF table and reachable from inet.0 by static configuration (for example, by configuring static/route/next-table/VRF.inet.0), and NTP source-address is configured, the ntpd (the Network Time Protocol daemon running on NTP client) might pick the wrong source-address instead the configured source-address. As a result, NTP server cannot reply the NTP packet back. [PR1150005](#)
- Two interrupts are received from the FPGA on the control board of the MX2010/MX2020 platforms for every i2c transaction triggered from software. Only one is expected. [PR1151674](#)
- During the unified ISSU upgrade, line cards may crash, causing service impact. When the line cards come up, there may be a programming issue as a secondary impact and some IFLs may not pass traffic. Affected line cards need to be rebooted to recover from this condition. [PR1152048](#)
- On MX Series routers with Junos OS Release 14.2R5-S1, when we specify a multiservice (ms-) interface to add a timestamp to Real-time Performance Monitor (RPM) probe

messages, it will cause the mspmand process to crash and the MS-MPC/MS-MIC keep crashing. As a workaround, we should configure RPM to perform timestamping either on the Routing Engine (Routing Engine-based RPM) or on an installed MPC Packet Forwarding Engine (Inline-RPM). [PR1152785](#)

- Fixed an issue with Inline Jflow where the Observation Domain field in exported IPFIX datagrams were always using the value attributed for LU0 in MPCs with multiple LUs per forwarding-engine. [PR1152854](#)
- On MX Series platform, when MPC goes down ungracefully, other MPC in the chassis will experience "destination timeout". Due to this event, auto fabric-healing will get triggered due to "destination timeout" condition. Due to the software issue the fabric-healing starts from Phase-1 and in some cases it can go upto Phase-2 causing all other MPCs to be restarted. [PR1156069](#)
- From Junos OS Release 15.1F5 and later, the hidden configuration statement "filter-list-template" will be enabled by default for all firewall filters on MX Series based platforms to use a common program on MX Series boards for all interfaces that use the same filter list. This can save MX Series board microkernel memory and DMEM memory. The hidden configuration statement "no-filter-list-template" can be configured to disable this behavior. [PR1157079](#)
- Fixed an issue where MX Series cards could crash while programming a firewall filter containing flexible-match-mask. [PR1157759](#)
- With Junos OS Release 15.1F2 and later, when inline sampling is enabled on MX Series-based FPC, the srdd (Sampling Route-Record Daemon) process would be created to maintain, collect, and export JFLOW records. On a regular time intervals, the srdd scans through the sampling database for any update/change in the record. In a scaled environment with more route churn, for example 1.14M routes, the scan process might hog CPU for more than 2.5 sec which leads to FPC crash. In some situations, the scan process can run for longer time without causing FPC crash, but it can cause BFD sessions to flap. [PR1158154](#)
- The following commit warning may be seen when using configure private and multi-line comments. This causes the commit to not complete: "warning: outgoing comment does not match patch". [PR1161566](#)

### ***Routing Policy and Firewall Filters***

- When a malformed prefix is used to test policy (command "test policy <policy name> <prefix>"), and the malformed prefix has a dot symbol in the mask field (e.g. x.x.x.x/24), the rpd process might crash. [PR1144161](#)
- From Junos OS Release 13.2R1, an attempt to commit a configuration with a dangling conditional policy referring a nonexistent/inactive routing-instance will be permitted. If we have a conditional policy referring an active routing-instance, deleting/deactivating this routing-instance and then committing will cause the rpd process crash. As a workaround, we should always make sure that conditional policies are referring active routing-instances. [PR1144766](#)

### ***Routing Protocols***

- On dual Routing Engine platform with GRES and NSR enabled, after Routing Engine switchover, the rpd might crash when trying to destroy a CNH NH (composite next hop, for example, it would be created in PIM, L3VPN, MVPN scenario and so on) with valid reference on it. It is because that during switchover (while backup rpd switches to master), there is a transition period where rpd switched to master mode but KRT is still in backup mode. If KRT (still in backup mode) receives a CNH addition followed by Route additions using this CNH during this phase, it would result in CNH in KRT with valid route references yet on expiry queue. It is hard to reproduce, in this case, it occurs after Routing Engine switchovers consecutively at two times. [PR1086019](#)
- IGMPv2 working in v2/v1 compatibility mode does not ignore v2 Leave messages received on a bridge-domain's L2 member interface. Moreover, an IGMP snooping membership entry for the respective group at this L2 member interface will be timed out immediately upon IGMPv2 Leave reception, even when there are some other active IGMP hosts attached to this L2 member interface. It might break multicast forwarding for this L2 member interface. [PR1112354](#)
- When two (or more) route target communities of MP-BGP route match to two (or more) route target communities in VRF import policy of a RI, duplicate routing entries might be installed in the RI. In the output of 'show route table <RI name>.inet.0 detail', two identical routing entries appear with one being marked as 'Inactive reason: Not Best in its group - No difference'. When such duplicate routing information is to be deleted, rpd process process will crash. [PR1113319](#)
- During many types of configuration changes, especially including import policy, BGP has the need to re-evaluate the routes it has learned from peers impacted by the configuration change. This re-evaluation involves re-running import policy to see if there is any changes to the learned routes after applying the new policy. This work is done in the background as part of an "Import Evaluation" job. When BGP is reconfigured a second time, and the "Import Evaluation job" has not completed, it is necessary to re-run the job from the beginning if there's another change to policy or something with similar impact. This state is noted as "Import Evaluation Pending". However, in this case, there was a bug that caused BGP to always enter the pending state upon reconfiguration, regardless of whether relevant changes were made to import or other similarly impactful configuration. The result is that once it is necessary to start re-evaluation of the routes for a peer, even trivial configuration changes that happen too quickly will cause the "Import Evaluation job" to need to run again as a result of the "Pending" flag being set. To avoid the issue, please ensuring that "ImportEval" is not present in a BGP peer's Flags output from the CLI (show bgp neighbor) prior to doing even trivial commits. [PR1120190](#)
- In multicast environment, when the RP is FHR (first-hop router) and it has MSDP peers, when the rpf interface on RP changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)
- On dual Routing Engine platform with Bidirectional Forwarding Detection (BFD) protocol enabled, after graceful Routing Engine switchover (GRES), the periodic packet

management process (ppmd) might crash on backup Routing Engine due to a software defect. [PR1138582](#)

- RPD cores while processing PIM hellos. There is no known workaround for this problem. RPD core seems to happen sometimes when a \*g and sg's vanish mostly due to LHR becoming a Non-DR from a DR. [PR1140230](#)
- With NSR configured, when the BFD sessions are replicated on backup Routing Engine, the master will not send the source address, instead backup Routing Engine will query the kernel to get the source address. In rare cases, the query might fail, resulting in the source address as all zeros. Later, if a GRES switchover happens, new master will have this all zeros source address. When BFD packet with this source address is sent out, the other end will drop the BFD session due to no matching session (source address). [PR1145612](#)
- Core seen when BMP station was passive, and the BMP Collector was terminated non-gracefully, and BMP station was not properly cleaned up. [PR1154017](#)
- When a BFD session is configured over an Aggregated Ethernet interface located on a MPC and the MX Series chassis is set to non-enhanced IP or Ethernet network service mode, with Junos OS Release 15.1F2 or later, the BFD session might be unstable. [PR1162716](#)

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

- Junoscript traceoptions are available. [PR1062421](#)
- When entering the "restart r" incomplete command in the CLI, the command "restart routing" is executed. It should throw an error like "error: invalid daemon: r". [PR1075746](#)
- From Junos OS Release 13.2R1 and later, the committed process might crash while committing large configurations in a single commit, for example, committing 250k lines of config on top of existing configuration. This issue is due to a lack of storage space for current and running configurations. [PR1159462](#)

#### ***VPNs***

- For a next-generation multicast VPN (NG-MVPN) using ingress replication provider tunnels, if both IPv4 and IPv6 are configured, when receiver PE advertises different labels for IPv4 and IPv6 in type-1 BGP route, the source PE will create two provider tunnels to carry IPv4 and IPv6 traffic both and causing duplicated multicast traffic. [PR1128376](#)
- If one VRF has Draft-Rosen 6 MVPN for IPv4 and Next-Generation MVPN for IPv6, when walking through SNMP MIB for MvpnSpmsiTable, the rpd process may hit NULL pointer and crash. The routing protocols are impacted and traffic disruption will be seen due to loss of routing information. [PR1145241](#)

#### ***Resolved Issues: 15.1F4***

- [Class of Service \(CoS\) on page 142](#)
- [Forwarding and Sampling on page 142](#)
- [General Routing on page 142](#)

- [Interfaces and Chassis on page 144](#)
- [Layer 2 Features on page 144](#)
- [MPLS on page 145](#)
- [Network Management and Monitoring on page 145](#)
- [Platform and Infrastructure on page 145](#)
- [Routing Protocols on page 146](#)
- [Services Applications on page 147](#)
- [VPNs on page 147](#)

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

- This PR does optimization in AE SNMP handling. If all the links in an AE bundle go down, then any COS SNMP query for this AE IFD/IFL will return cached values. [PR1140440](#)

### ***Forwarding and Sampling***

- On all Junos OS platforms, when both the filter and the policer are configured for an interface, in rare cases, the policer template may not be received by Packet Forwarding Engine (from the Routing Engine) when it is referenced by the filter term (normally the policer template gets received before the filter term referencing it which is ensured by mechanism in Routing Engine kernel). In this situation, the FPC would crash due to this timing issue. This issue might be avoid by the recommended steps below: 1. Deactivate the physical interface (IFD) and commit 2. Enable any filter and policer that attached to the interface (e.g. IFL) and commit 3. Activate interface back. [PR1128518](#)

### ***General Routing***

- No performance or functional impact. Can be safely ignored. "Ignore the PTP message (2) as this MPC does not support EEC" should be moved from notice to debug level. [PR1020161](#)
- There is a remote loopback feature in 802.3ah standard, where one end can put the remote end into remote-loopback mode by sending an enable loopback control LFM PDU. In remote loopback, all incoming packets (except LFM packets) are sent back on wire as it is. Transmit or receive of LFM packets should not be affected when an interface is in remote loopback mode. On VMX platforms, when we configure the LFM remote-loopback we run into problem state. In problem state we will see that LFM packets sent from node which is in loopback state is not reaching the peer end, hence we will not see the remote entity information for the "run show oam ethernet link-fault-management" command on the peer router. [PR1046423](#)
- After executing CLI command "show route extensive", routing protocol process (rpd) may get into infinite loop and not respond anymore because the command may get executed a couple of times itself. In this situation, rpd high CPU utilization (running over 90% sometimes) might be seen on the device, and also the memory which used to store the command output would not be freed during those executions (in normal utilization, the memory uses about 160KB, but in problematic situation, it can swell to 3GB size), which would lead to rpd crash eventually after memory exhaustion. [PR1104090](#)

- When Bridge domain in PBB-EVPN Routing instance is modified to add/remove ISIDs BD can get stuck in destroyed state. This happens when ISIDs in the Bridge domain are changed from 1 to many or many to 1. This is only noticed during configuration changes or initial deployment. [PR1107625](#)
- In rare condition, after Routing Engine switchover, the MPC PIC might be offline, and some error messages might be seen. [PR1110590](#)
- On dual Routing Engine MX Series platform, the "xe" interfaces of any of the line cards below may flap during unified in-service software upgrade (ISSU) due to missing support. The flapping may not happen every time and the probability of occurrence would increase if more number of SFP+ (e.g., SFP+-10G-SR) are connected on the FPC. The affected line cards are, \* MIC3-3D-10XGE-SFPP \* MPC4E-3D-32XGE-SFPP, MPC4E-3D-2CGE-8XGE \* MPC5E-40G10G, MPC5EQ-40G10G \* MX2K-MIC6-24XE, MX2K-MIC6-24XE-OTN. [PR1118379](#)
- On MX240/MX480/MX960/MX2010/MX2020 products with MPC2E-3D-NG/MPC2E-3D-NG-Q/MPC3E-3D-NG/MPC3E-3D-NG-Q with MIC-3D-4XGE-XFP, IFD flap detection is much slower. It might lead to high FRR time, some traffic might be lost. [PR1122589](#)
- With BGP configured on CE-faced interfaces (in VRFs), doing 'show route' frequently may cause rpd to slowly leak memory. The leak rate will be one memory block of the size necessary to hold the instance name of the routing instance for a BGP neighbor. If the rpd process memory is exhausted, the rpd process might crash, and the routing protocols are impacted and traffic disruption will be seen due to loss of routing information. You can check rpd memory usage with "show task memory brief" command. [PR1124923](#)
- In multihoming EVPN scenario and the customer facing interface is an AE interface, after moving an interface from the EVPN instance into a VPLS instance, traffic loss might be seen on CE facing FPC. [PR1126155](#)
- EVPN route attributes like the label and Ethernet segment identifier (ESI) may be missing from EVPN family routes installed by BGP. [PR1126770](#)
- In 15.1F3, rpd core can be seen on previous master after performing Routing Engine switchover. [PR1128023](#)
- In current Juniper Networks implementation, the IPv6 multicast Router Advertisement timer is not uniformly distributed value between MinRtrAdvInterval and MaxRtrAdvInterval as described in RFC 4861. [PR1130329](#)
- On MX Series based line card, multiple modifications of firewall filter might cause lookup chip error and traffic blackhole, following jnh\_free error messages could help to identify this issue: "messages: fpc1 jnh\_free(10212): ERROR [FW/3]:1 Paddr 0x006566a9, addr 0x2566a9, part\_type 0 call\_stack 0x40497574 0x418ffa84 0x41900028 0x418ecf94 0x41861690". [PR1131828](#)
- 100G interface in MPC3E is not coming up after unified ISSU in sync. [PR1136269](#)

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

- The adaptive load balancing counters are always zero for aggregated Ethernet (AE) bundles on MICs or MPCs of MX Series routers. [PR1101257](#)
- The following CLI configuration statement needs to be used for CFM session to work. "set chassis aggregated-devices disable-lag-enhanced". Enhanced-lag is enabled by default in the system when the system is configured with enhanced-ip. CFM is not supported with enhanced-lag at present. [PR1116826](#)
- On Junos OS platform, an aggregate-ethernet bundle having more than one member link can show incorrect speed which would not match to the total aggregate bandwidth of all member links. The issue would be seen when LFM is enabled on the aggregate-ethernet bundle. The issue would be triggered when one of the member link flaps. Although after the flap, the current master Routing Engine would show correct aggregate speed, the backup Routing Engine would report incorrect value. In this state, when Routing Engine mastership is switched, the new master Routing Engine (which was backup) will show incorrect value. One of the side-effect of this issue is that RSVP also reflects incorrect Bandwidth availability for the affected aggregate-ethernet bundle, thus can cause under-utilization of the link with LSP having bandwidth constraints. [PR1121631](#)
- Since a bug which was introduced in 15.1R1, loopback sub-interfaces always have a Flag down in the output of CLI command "show interfaces". [PR1123618](#)
- The connectivity fault management (CFM) log message "Adjacency up" should only be logged when the router first detects remote MEP or the peer interface goes down and up causing adjacency failure for this remote MEP. But now it is incorrectly logged when any peer set/clear the Remote defect indication (RDI) bit in continuity check messages (CCMs). [PR1125164](#)

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

- For Routing Engine generated packet with VLAN tag, if the outgoing interface is an LT interface, the VLAN tag will not be removed even the LT interface is configured with untagged encapsulation. [PR1118540](#)
- In some rare scenarios, the MVRP PDU might be unable to be transmitted, which could cause memory leak in layer 2 control plane daemon (l2cpd), and finally results in the l2cpd process crash. [PR1127146](#)



### **MPLS**

- When local bandwidth accounting for inactive /adaptive standby path figures that there is not enough bandwidth to fit it in an already full link and brings it down, CSPF will not be retried on the path unless there is some change in TE database. [PR1129602](#)

### **Network Management and Monitoring**

- On Junos OS releases 13.1X42/14.1X51/15.1R1/15.1R2, the SNMP average response time in the output of "show snmp statistics extensive" is incorrectly calculated and might be observed with negative value. [PR1112521](#)

### **Platform and Infrastructure**

- When one of the "deny-commands" is incorrectly defined on the profile of TACACS+ server, all "deny-commands" regexes are ignored, which leads to an over-permissive profile without any warning. [PR1078238](#)
- In 64-bit Junos OS environment, the Representational State Transfer (REST) API fails to start when configured with "set system services rest ...". [PR1097266](#)
- After changing an outer vlan-tags, the ifl is getting programmed with incorrect stp state (discarding), so the traffic is getting dropped. [PR1121564](#)
- When MX2020 or MX2010 is running with FreeBSD10-based 15.1 Junos OS image, I2C error will be seen sporadically. tcbc i2c accelerator error: Group 0xX device 0xXX cmd timedout 984 usecs If the i2c error happens on voltage sensor, and it reaches count limit (9 times), chassis alarm will be shown up like this. 1 alarms currently active Alarm time Class Description 2015-09-10 06:42:40 UTC Minor CB 1 Volt Sensor Fail Those are cosmetic error but there is no way to clear the chassis alarm other than offline/online the FRU. [PR1122821](#)
- On MX Series-based platform, when fragmented packets go through the inline NAT (including source NAT, destination NAT, and twice NAT), the TCP/UDP checksum would not be correctly updated. In this situation, checksum error would occur on the remote end (inside and outside device). Non-fragmented packets would not be affected by the issue. If possible, this issue could be avoided by either of the following workarounds: \* Enable "ignore-TCP/UDP-Checksum errors" at the inside or outside device which processes TCP/UDP data, OR \* Make sure there will not be any fragments subjected to inline NAT functionality by appropriate MTU adjustment or setting. [PR1128671](#)
- Parity error at ucode location which has instruction init\_xtxn\_fields\_drop\_or\_clip will lead to a LU Wedge. LU is lookup ASIC inside the MX Series platform. The LU wedge will cause the fabric self ping to fail, which will lead to a FPC reset. This is a transient HW fault, which will be repaired after the FPC reset. There is no RMA needed unless the same location continues to fail multiple times. [PR1129500](#)
- NTP.org published a security advisory for thirteen vulnerabilities in NTP software on Oct 21st, 2015. These vulnerabilities may allow remote unauthenticated attackers to cause Denial(s) of Service(s), disruption of service(s) by modification of time stamps being issued by the NTP server from malicious NTP crafted packets, including maliciously crafted NTP authentication packets and disclosure of information. This

can impact DNS services, as well as certificate chains, such as those used in SSL/https communications and allow attackers to maliciously inject invalid certificates as valid which clients would accept as valid. [PR1132181](#)

- PPE thread timeout trap may cause XM chip wedge; it will not affect MQ-based FPC. [PR1136973](#)
- On MX2020, when we remove whole power of a power zone, and then put the power back to the zone, FANTray LED stays Amber and FANTray LED on craft card stays OFF, and do not revert to green (FANTray LED) or ON (Craft LED) until we reboot the entire chassis system or hot swap that FAN tray. For Zone 0 (PSM 0 to 8), FAN 1 shows the above described behavior. For Zone 1 (PSM 9 to 17), FAN 3 shows the above described behavior. [PR1138209](#)

### ***Routing Protocols***

- There may be stale BFD session after changing physical interface MTU. It may also cause BFD session to flap continuously or to stay in down state. [PR1116666](#)
- When an interface is associated with a Bidirectional Forwarding Detection (BFD) session, if changing the unit number of the interface (for example, change the unit number for a running BFD session from ge-1/0/0.2071 to ge-1/0/0.285), the device may fail to change the name due the missing check for logical interface (IFL) index change. In addition, this is a software issue and may not have any service impact. [PR1118002](#)
- When protocol MSDP is configured and then deleted, the NSR sync status for MSDP might stuck in "NotStarted", and unified ISSU might fail on master Routing Engine with reason "CHASSISD\_ISSU\_ERROR: Daemon ISSU Abort -1(NSR sync not complete: MSDP)". [PR1129003](#)
- In multicast environment, when the RP is FHR (first-hop router) and it has MSDP peers, when the rpf interface on RP is changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)
- On dual Routing Engine platforms, due to software issue, OSPF (including both OSPFv2 and OSPFv3) "DoNotAge" bit (e.g., source of LSA has flood-reduction feature enabled) is not mirrored to backup routing protocol process (rpd). In this situation, after performing nonstop active routing (NSR) switchover, the LSA on new master rpd remains without "DoNotAge" bit set. Once the LSA reaches OSPF max age, the router will flood LSA purge, hence route flapping might be seen on all routers under the OSPF topology. [PR1131075](#)
- When applying add-path prefix-policy to neighbor level, all neighbors are separated into different update groups. This is not the expected behavior. There is no service impact. But if all the neighbors are configured under one peer group with huge number of peer groups, the scaling/performance will go down. [PR1137501](#)

### Services Applications

- The Point-to-Point Tunneling Protocol (PPTP) ALG is used for tunneling PPP packets over an IP network. But if the router configures session-limit-per-prefix, the PPTP-ALG does not work. [PR1128484](#)

### VPNs

- In L2circuit environment, if one PE has pseudowire-status-tlv configured but remote has not, and at the same time, this PE does not support control-word but remote does, then it will not send changed local status code to remote PE. In a rare condition, after enable status-tlv support at remote end, the l2circuit might get stuck in "RD" state on remote PE. [PR1125438](#)

### Resolved Issues: 15.1F3

---

- [General Routing on page 147](#)
- [Infrastructure on page 150](#)
- [Interfaces and Chassis on page 151](#)
- [Layer 2 Features on page 152](#)
- [MPLS on page 152](#)
- [Network Management and Monitoring on page 153](#)
- [Platform and Infrastructure on page 153](#)
- [Routing Protocols on page 155](#)
- [Software Installation and Upgrade on page 156](#)
- [VPNs on page 156](#)

### General Routing

- In MX Virtual Chassis (MX-VC) environment, if the private local next-hops and routes pointing to private local next hops are sent to Packet Forwarding Engine from the master Routing Engine and not sent to the backup Routing Engine, then a Routing Engine switchover happens. Now as the new master Routing Engine does not know about such next hops and routes, they are not cleaned up. When a next hop with same index is added on new master Routing Engine and sent to Packet Forwarding Engine, the Packet Forwarding Engine might crash due to a stale next hop. [PR951420](#)
- In a Layer 3 wholesale configuration, DHCPv6 advertise messages might be sent out with source MAC all zeros if the subscriber is terminated on the demux interface in a non-default routing instance. For subscribers on default instance, there is no such issue observed. [PR972603](#)
- Earlier the output of "show agent sensors | display xml" used to show sensor details and the attached server and export-profile details at the same level in xml output. This is confusing since there are multiple sensor data listed for this command and all will be shown with same indentation. After this change, the output of "show agent sensors | display xml" will be shown as the following with each <sensor> tag covering a single sensor's xml data: root@Router# run show agent sensors | display xml <rpc-reply xmlns:junos=URL>

```
<sensor-information>
<sensor>
<sensor-name>name-of-sensor-here </sensor-name>
<resource-name> resource-path </resource-name>
<sensor-id>scope-id</sensor-id>
<resource-filter>resource-filter-name </resource-filter>
<server-information>
<server-name>streaming-server-name </server-name>
<scope-id>scope-id</scope-id>
<remote address>remote-address </remote address>
<remote-port>remote-port</remote-port>
</server-information>
<profile-information>
<profile-name>export-profile-name </profile-name>
<rep-interval>reporting-interval</rep-interval>
<local-address>local-address </local-address>
<local-port>local-port </local-port>
<timestamp>timeticks </timestamp>
<serverformat>server-export-format </serverformat>
<transport>transport </transport>
<dscp>code-point </dscp>
<forwarding-class>forwarding-class </forwarding-class>
</<profile-information>
</sensor>
<sensor>
...
</sensor>
```

#### [PR1037064](#)

- Upon BFD flapping on aggregate interfaces, the Lookup chip (XL) might send illegal packets to the center chip (XMCHIP) and compromise packet forwarding and an FPC restart is needed to recover from this condition. If Fabric path side is affected, the fabric healing process will initiate this process automatically to recover from such conditions. Only MPC5E or MPC6E are exposed to this problem. [PR1067234](#)
- ICMP echo\_reply traffic with applications like IPsec will not work with the MS-MIC and MS-MPC cards in an asymmetric traffic environment since these cards employ a stateful

firewall by default. The packet will be dropped at the Stateful Firewall since it sees an ICMP Reply that has no matching session. [PR1072180](#)

- Remnant routes seen in old master Routing Engine after Routing Engine switchover in non-GRES scenario. [PR1075404](#)
- In a two member MX Series Virtual Chassis (MXVC) environment, when "set virtual-chassis no-split-detection" is configured, if split master condition happens, which is caused by split events (i.e., loss of all adjacencies by link failure, FPC restarts, chassis power-down, Routing Engine reboots, etc), then once the VCP adjacency is formed again, the current design could not determine the best chassis to win the protocol mastership election properly. Instead, only the final election step (that is, choose the member device with the lowest MAC address) is used to elect the master device (protocol master of the VC, or VC-M). [PR1090388](#)
- The OpenSSL project has published a set of security advisories for vulnerabilities resolved in the OpenSSL library in June and July 2015. Junos OS is affected by one or more of these vulnerabilities. Refer to JSA10694 for more information. [PR1095598](#)
- When route convergence occurred, the new gateway address is not updated correctly in inline-jflow route-record table (route-record table is used by sampling), and the sampling traffic forwarding might be affected, but normal routing would be not affected. [PR1097408](#)
- Some of the new revisions (for example, REV 30, REV 31) of the MICs cannot come up with NG-MPC2 or NG-MPC3 line card. We can check the MIC version by CLI command "show chassis hardware detail | no-more".  

```
root@user> show chassis hardware detail | no-more
Hardware inventory:
Item Version Part number Serial number Description ..
FPC 2 REV 14 750-054901 CADJ3871 MPC3E NG PQ & Flex Q CPU REV 11 711-045719
CADN5465 RMPC PMB MIC 0 REV 30 750-028392 CAEB9203 3D 20x 1GE(LAN) SFP
<<<<<<REV>PR1100073
```
- After Junos OS Release 13.3R1, IPCMON infrastructure is added to debug IPCs between PFEMAN and the Routing Engine. When convergence occurs, string processing of IPCMOM will take added time. Then the slow convergence will be seen. It is a performance issue, and it is visible in scaled scenario (for example, more than 100K routes). As a workaround, execute the command "set pfe ipclog filter clear" to disable IPC logging on all FPCs. [PR1100851](#)
- In broadband edge (BBE) environments, for example, if the interface-set is created corresponding to SVLAN, then multiple logouts and logins will create a new interface-set index. When the interface-set index range goes above 65535, executing CLI command "show interfaces interface-set queue egress" will cause 100% CPU usage. As a workaround, we can use the specified interface-set name instead of using the wildcard. [PR1101648](#)
- Non-queuing MPC5E might crash continuously if rate-limit under transmit-rate for scheduler is applied. As a workaround, do not configure rate-limit and use firewall policer for forwarding-class instead. MPC5EQ is not exposed. [PR1104495](#)
- A vulnerability in OpenSSH may allow a remote network based attacker to effectively bypass restrictions on number of authentication attempts, as defined by MaxAuthTries settings on Junos OS. This may enable brute force password attacks to gain access to

the device. Background: The PAM (Pluggable Authentication Modules) library provides a flexible framework for user authentication and session setup / teardown. It is used not only in the base system, but also by a large number of third-party applications. Various authentication methods (UNIX, LDAP, Kerberos etc.) are implemented in modules which are loaded and executed according to predefined, named policies. These policies are defined in `/etc/pam.conf`, `/etc/pam.d/<policy name>`, `/usr/local/etc/pam.conf`, or `/usr/local/etc/pam.d/<policy name>`. The PAM API is a de facto industry standard which has been implemented by several parties. FreeBSD uses the OpenPAM implementation. This issue is assigned CVE-2015-5600. [PR1106752](#)

- On MX Series platforms with "subscriber-management" enabled, while high-scaled subscribers (for example, 126K dual-stack DHCP v4/v6 subscribers over VLAN demux) log in/log out at high rate, MPCs and MICs that hold subscribers might crash after the bbe-smgd process restarts. [PR1109280](#)
- In the scenario that the power gets removed from the MS-MPC, but the Routing Engine is still online (for example, on MX960 platform with high-capacity power supplies that split into two separate power zones, when the power zone for the MS-MPC line card loses power by switching off the PEM that supports the MS-MPC situated slot), if the power goes back on (for example, switch on the PEM), the MS-MPC might be seen as "Unresponsive" (checked via CLI command "show chassis fpc") and not coming up back online due to failure of reading memory. [PR1112716](#)
- On MX-VC with heartbeat connection, if it is in a scaled subscribers environment, when power down both VCM Routing Engines, there might be a delay (minutes) for backup chassis to be master and during which time, traffic blackhole might be seen. [PR1115026](#)
- No decrement ttl does not work for incoming v6 traffic over MPLS IPv4 core. [PR1115203](#)
- For MPC6E with CFP2, there was a race condition between the Interrupt service routine and the periodic, as a result interface up/down will not happen for laser off/on. [PR1115989](#)
- On MX Series platforms, the 10G Tunable SFP/SFP+ cannot be tuned in Junos OS Release 15.1R2. [PR1117242](#)
- The rpd process might crash when executing CLI command "show evpn database" with the combination of "vlan-id" and "mac-address". [PR1119301](#)

### Infrastructure

- Only the following directories and files are preserved when upgrading from a build prior to Release 15.1 to Release 15.1 (FreeBSD 10): `config/` `/etc/localtime` `/var/db/` `/var/etc/master.passwd` `/var/etc/inetd.conf` `/var/etc/pam.conf` `/var/etc/resolv.conf` `/var/etc/syslog.conf` `/var/etc/localtime` `/var/etc/exports` `/var/etc/extensions.allow` `/var/preserve/` `/var/tmp/baseline-config.conf` `/var/tmp/preinstall_boot_loader.conf`. Anything else not listed above is deleted/formatted during the upgrade to the freebsd10 version of Junos OS. [PR959012](#)
- When "show version detail" CLI command has been executed, it will call a separate gstatd process with parameter "-vvX". Because the gstatd could not recognize these parameters, it will run once without any parameter then exit. In result of "show version detail", following information could be seen: `user@mx960> show version detail`

Hostname: mx960 Model: mx960 Junos: 13.3R6-S3 JUNOS Base OS boot [13.3R6-S3] JUNOS Base OS Software Suite [13.3R6-S3] JUNOS Kernel Software Suite [13.3R6-S3] JUNOS Crypto Software Suite [13.3R6-S3] <snipped> file: illegal option -- v usage: gstatd [-N] gstatd: illegal option -- v usage: gstatd [-N] <snipped> At the same time, log lines like following might be recorded in syslog: Aug 25 17:43:35 mx960 file: gstatd is starting. Aug 25 17:43:35 mx960 file: re-initialising gstatd Aug 25 17:43:35 mx960 mgd[14304]: UI\_CHILD\_START: Starting child '/usr/sbin/gstatd' Aug 25 17:43:35 mx960 gstatd: gstatd is starting. Aug 25 17:43:35 mx960 gstatd: re-initialising gstatd Aug 25 17:43:35 mx960 gstatd: Monitoring ad2 Aug 25 17:43:35 mx960 gstatd: switchover enabled Aug 25 17:43:35 mx960 gstatd: read threshold = 1000.00 Aug 25 17:43:35 mx960 gstatd: write threshold = 1000.00 Aug 25 17:43:35 mx960 gstatd: sampling interval = 1 Aug 25 17:43:35 mx960 gstatd: averaged over = 30 Aug 25 17:43:35 mx960 mgd[14304]: UI\_CHILD\_STATUS: Cleanup child '/usr/sbin/gstatd', PID 14363, status 0x4000 Aug 25 17:43:35 mx960 mgd[14304]: UI\_CHILD\_EXITED: Child exited: PID 14363, status 64, command '/usr/sbin/gstatd' [PR1078702](#)

- On MX Series platforms with Junos OS Release 15.1R1 or above, while a core dump is in progress, if we try to access the dump directory, due to the deadlock defect, the system might hang and crash. As a workaround, we should not access the "/var/crash" directory till the core dump is complete. [PR1087082](#)
- On dual Routing Engine platform, if GRES is configured (triggered by "on-disk-failure"), when a disk I/O failure occurs on the master Routing Engine due to hardware issue (for example, SSD failure), the graceful Routing Engine switchover might not be triggered immediately after initial IO failure has been detected. As a result, Routing Engine might enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. [PR1102978](#)
- With scaled configuration or there are memory leaks, if the virtual memory is running very low, the kernel might crash and the device will go in db prompt continuously due to a recursion issue. [PR1117548](#)
- show route vpn-localization command does not show any output, but if xml format is requested, then xml output of the same command works. [PR1125280](#)

### **Interfaces and Chassis**

- On MX Series router, the physical or logical interfaces (ifd/ifl) might be created and marked UP before resetting FPCs' fabric planes are brought up and ready to forward traffic. As a result, traffic might be black-holed during the time window. This window of traffic black-hole is particular long if the chassis is heavily populated with line-cards, for example, the router has large scale of configuration (routes or subscribers), and coupled with a lot of FPC resets, such as upon a node power up/reset. [PR918324](#)
- When issuing a CFM LTR from CE, link state reply, received from MX Series, acting as MHF does not contain Reply Egress TLV if ingress and ingress logical interface are located on the same IFD. [PR1044589](#)
- During subscriber login/logout, the following error log might occur on the device configured with GRES/NSR: /kernel: if\_process\_obj\_index: Zero length TLV! /kernel: if\_pfe: Zero length TLV (pp0.1073751222). [PR1058958](#)

- For Junos OS Release 13.3R1 and later releases, after multiple (e.g., 26) iterations of graceful Routing Engine switchover (GRES), the TNP address of management interface might be deleted incorrectly during switchover, which leads to all FPCs to be offline. [PR1060764](#)
- Trap messages do not get logged on logical interface (ifl) after deleting "no-traps" configuration statement, in spite of setting explicit "traps". [PR1087913](#)
- On MX240 or MX480 platform with at least two DC modules (PN: 740-027736) equipped, when shutting down one of the PEMs and then turning it on again, even when the PEM is functioning, the "PEM Fan Fail" alarm might be observed on the device due to software logic bug. There is no way to clear the ALARM\_REASON\_PS\_FAN\_FAIL for I2C\_ID\_ENH\_CALYPSO\_DC\_PEM once it has been raised. [PR1106998](#)
- On all Junos OS platforms, if the "HDD /var" slice (for example, "/dev/ad1s1f" depending on the type of Routing Engine) is not mounted (for example, label missing, file system corrupted beyond repair, HDD/SDD is removed from the boot list, etc), the system may build emergency "/var/", however, no alarm or trap is generated due to the incorrect operation of the ata-controller. Although the boot messages may present the logs, it may not be sufficient enough to identify the issue before encountering other problems (for example, Junos OS upgrade failure and the Routing Engine may hang in a recovery shell). In addition, as a method to check where Routing Engine is running from, a manual check could be done as below, user@re0> show system storage | match " /var\$"  
/dev/ad2s1f 34G 18G 13G 57% /var <<<<Indicate that>show system storage | match " /var\$" <<<<NO output> [PR112580](#)

### **Layer 2 Features**

- With Dynamic Host Configuration Protocol (DHCP) maintain subscriber feature enabled, when the subscriber's incoming interface index is changed, for example, the interfaces go away and come back after changing the MTU configuration of interface, the existing subscribers may get dropped and new subscribers fail in connection. [PR1059999](#)
- For PVSTP/VSTP protocols, when MX Series router inter-operates with Cisco device, due to the incompatible BPDU format (there are additional 8 bytes after the required PVID TLV in the BPDU for Cisco device), the MX Series router might drop these BPDUs. [PR1120688](#)

### **MPLS**

- On dual Routing Engine platform with GRES, the kernel synchronization process (ksyncd) may crash on the backup Routing Engine when adding of route pointing to indirect nexthop on system. [PR1102724](#)
- From Junos OS Release 13.2R1 and later, in MPLS L3VPN scenario, when "l3vpn-composite-nexthop" configuration statement is enabled on a PE router and an interface style service set is attached to the ingress interface, the L3VPN packets with the MPLS labels will be sent to the service card and dropped. As a workaround, we should disable "l3vpn-composite-nexthop". [PR1109948](#)
- If "optimize-timer" is configured under P2MP branch LSP, this branch LSP will not be re-established if link flap on egress node. If "optimize-timer" is configured at protocols/mpls level, issue could be avoided. [PR1113634](#)



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

- In rare cases, when the mib2d process attempts connection with the snmpd process and there are pending requests waiting to be finished, the mib2d process might crash and the CPU utilization is high around the same time as the crash happens. [PR1076643](#)
- While the router is rebooting and SNMP polling is not stopped, SNMP requests might land on mib2d process before Routing Engine protocol mastership is resolved, causing the mib2d process crash. [PR1114001](#)

### ***Platform and Infrastructure***

- 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 fails. [PR872609](#)
- On MX Series based line card, when GRE keepalive packets are received on a Packet Forwarding Engine that is different from the tunnel interface hosted, the keepalive message will apply the firewall filter configured on default instance loopback interface. [PR934654](#)
- When the 'enhanced-hash-key services-loadbalancing' feature is used by MX Series based line cards, load balancing flows across multiple service PICs via the source-address do not work when internal BGP (IBGP) is used to steer traffic to the inside service-interface. For example, the operator will see on the stateful firewall that the same source-address has flows across multiple service interfaces. [PR1034770](#)
- If with both MPC/MSDPC and other type of DPCs equipped, for local switching at mesh group level, split horizon on PW interfaces will not work and this would cause packets to loop back to same PW interface. [PR1084130](#)
- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)
- On MX Series platform, if ingress "multicast-replication" is configured, the throughput of the multicast may get reduced due to unnecessary threads during Packet Forwarding Engine operation. In addition, only the performance of multicast traffic may get influenced (some of the multicast packets may get dropped on the Packet Forwarding Engine) by the issue. This PR has fixed/enhanced the performance. Now the performance limit should only be capped by fabric bandwidth in ingress Packet Forwarding Engine. In addition, before this fix, there was a limitation that VPLS/Bridging can't run with ingress-replication feature as its BUM traffic can't be handled by ingress-replication feature. This PR removed that limitation as well. Now BUM traffic for VPLS/Bridging is following normal multicast replication path even with ingress-replication feature. [PR1089489](#)
- In 64-bit Junos OS environment, the Representational State Transfer (REST) API fails to start when configured with "set system services rest ...". [PR1097266](#)
- On MX Series platform, if ingress "multicast-replication" is configured, the throughput of the multicast may get reduced due to unnecessary threads during Packet Forwarding Engine operation. In addition, only the performance of multicast traffic may get

influenced (some of the multicast packets may get dropped on the Packet Forwarding Engine) by the issue. [PR1098489](#)

- On MX Series-based platform, before creating a new unicast next hop, there is a check to see if there are at least 512k DoubleWords (DWs) free. So, even the attempting next hop requires only a small amount of memory (for example, < 100 DWs). If there is not enough free DWs (that is, 512k), the check will fail, and the end result is that the control plane will quit adding this next hop prematurely - stopping at ~80% of capacity. With the fix, it will check for 64k free DWs which is a lower reference watermark for the available resource, thereby ensuring that it can allocate resource. [PR1099753](#)
- From Junos OS Release 14.1 and later, IPv6 mobility packets with Heartbeat option that the length of the mobility header (including the Ethernet encapsulation and main IPv6 header) extends beyond 128 bytes will be discarded as bad IPv6 option packet due to a logic error in packet handling. [PR1100442](#)
- Large scaled inline BFD session (in this case, 6000 inline BFD sessions) are loaded with the minimum-interval value 50ms. If FPC restarts, some BFD sessions might flap. [PR1102116](#)
- A remote attacker can cause a denial of service to the MX Series Chipset (Trinity) MPC due to maliciously crafted uBFD packets that are received directly, via VPN, MPLS, multicast, broadcast, on vt-interfaces, or otherwise. This issue affects both IPv4 and IPv6 traffic in both ethernet, and non-ethernet physical environments, such as ATM, or SONET, where the crafted packet is received over physical interfaces. If processed from a DPC through to the MPC then in-transit traffic will not be susceptible. In 6PE scenario, if the system is not using LSI/vt then not susceptible. If processed via MPC line card will be affected, the MPC line card will crash. If processed via endpoint receiving MPC line card terminating tunneling protocols such as MPLS/IPSec VPNs, etc. will be affected, this is considered in-transit traffic scenario. This crash can happen when the crafted packet is directed directly to the lo0 interface IP/physical interface IP/broadcast IPv4 / IPv6 address of the Physical interface As a workaround, we can apply a control plane (lo0) filter to drop uBFD packets. This issue is assigned CVE-2015-7748. More detailed information in the below link:  
[https://kb.juniper.net/InfoCenter/index?page=content&id=JSA10701&cat=SIRT\\_1&actp=LISTPR1102581](https://kb.juniper.net/InfoCenter/index?page=content&id=JSA10701&cat=SIRT_1&actp=LISTPR1102581)
- On MPC3E/MPC4E line card, when the feature "flow-detection" is enabled (under "ddos-protection" hierarchy), if suspicious control flow is received, two issues may occur on the device: Issue 1: sometimes, the suspicious control flow may not get detected on the line cards. Issue 2: once the suspicious control flows are detected, they may never time out even if the corresponding packets stop. [PR1102997](#)
- The following fields have been added to v10 Sampling (IPFIX) template and data packets: - SAMPLING RATE - SAMPLING INACTIVE TIMEOUT - SAMPLING ACTIVE TIMEOUT - TOTAL PACKETS EXPORTED - TOTAL FLOWS EXPORTED. [PR1103251](#)
- On MX Series platform, when using the 64-bit image, if the configuration statement "source-address" is configured for the "radius-server" as the following, the RADIUS request may not be sent to RADIUS server due to the failure of setting the "source-address" on the device. `user@re0> show configuration system radius-server .. source-address 10.1.1.1; <<<<< The configuration statement that may cause the issue.` [PR1103517](#)

- Due to a software defect found in Release 13.3R7.3 and Release 14.1R5.4 inclusively, Juniper Networks strongly discourages the use of Junos OS 13.3R7.3 on routers with MQ-based MPC. This includes MX Series with MPC1, MPC2; and all mid-range MX Series. [PR1108826](#)
- DHCP End options (option 255) is missing by DHCP-relay agent (where 20 bytes DHCP options 82 inserted) for client DHCP discover message with 19bytes padding. [PR1110939](#)
- An IPv4 filter configured to use the filter block with term that has both "from precedence" and another non 5-tuple (i.e., not port, protocol, address) will cause an XL/EA based board to reboot. Example: set firewall family inet filter FILTER fast-filter-lookup set firewall family inet filter FILTER term TERM from precedence PRECEDENCE set firewall family inet filter FILTER term TERM from tcp-established. [PR1112047](#)
- MXVC- Traffic being dropped on egress VCP Packet Forwarding Engine (invalid fabric token) [PR1112752](#)
- When inline BFD sessions and inline jflow are configured on the same Packet Forwarding Engine, with the increasing of active flows (about 65k), the BFD session might flap constantly and randomly because the outgoing BFD packets are dropped. [PR1116886](#)
- Inline 6rd and 6to4 support for XL and XL-XM based platforms. [PR1116924](#)
- On MX Series-based FPC, when MPLS-labeled fragmented IPv6 packets are arriving at PE router (usually seen in 6PE and 6VPE scenario), the Packet Forwarding Engine might mistakenly detect such IPv6 header and then drop these packets as "L3 incompletes" in the output of "show interface extensive". [PR1117064](#)
- When static inline NAT translation is used, if the translated source-prefix or destination-prefix is modified for one NAT rule, it may impact the other NAT rules as well. [PR1117197](#)
- On MX Series-based line card, the firewall filter may have some issues when matching on Authentication Header (AH) protocol. This can affect VRRP (among others) when authentication is used, and a Routing Engine firewall filter is matching on protocol AH. As a workaround, we can change the filter to match on other criteria (e.g., source or destination address). [PR1118824](#)

### ***Routing Protocols***

- Issue in populating isisRouterTable values. Some entries are not filled correctly. This does not block/affect the functionality of IS-IS or other components. [PR1040234](#)
- On large-scale BGP RIB, advertised-prefixes counter might show incorrect value due to a timing issue. [PR1084125](#)
- When polling SNMP OID isisPacketCounterTable 1.3.6.1.2.1.138.1.5.3, the rpd process might crash. [PR1101080](#)
- Static BFD does not update interface name after changing the interface unit name. [PR1118002](#)

### **Software Installation and Upgrade**

- In certain conditions, when /var is not mounted from a persistent filesystem, executing a Junos OS upgrade will have unexpected results. This is caused by an inexact check of whether we are running from an Emergency VAR. [PR1112334](#)

### **VPNs**

- In scenario involving pseudowire redundancy where CE facing interface in the backup neighbor (can be non-standby, standby, hot-standby type), if the virtual circuit (VC) is not present for the CE facing interface, the CE facing interface may go up after committing an unrelated VC interface configuration (e.g., changing description of another VC interface) even though the local pseudowire status is in down state. [PR1101886](#)

---

### **Resolved Issues: 15.1F2**

- [Class of Service \(CoS\) on page 156](#)
- [General Routing on page 157](#)
- [Infrastructure on page 159](#)
- [Interfaces and Chassis on page 159](#)
- [Layer 2 Features on page 161](#)
- [MPLS on page 161](#)
- [Network Management and Monitoring on page 162](#)
- [Platform and Infrastructure on page 162](#)
- [Routing Protocols on page 163](#)
- [Services Applications on page 164](#)
- [Software Installation and Upgrade on page 165](#)
- [User Interface and Configuration on page 165](#)
- [VPNs on page 165](#)

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

- In SNMP environment, when performing multiple walks or parallel snmpget for same interface at the same time (for example, SNMP bulk get/walk, or SNMP polling from multiple devices) on CoS related MIBs (jnxCos table), if the interface state changes or the request gets timeout when FPC is responding the request, memory leak of Class-of-Service process (cosd) about 160 bytes (up to 1500 bytes) may occur, which may cause cosd to crash eventually when limit is exceeded. [PR1058915](#)
- On MX Series platforms configured for IP network-services (default) and with MS-DPC/Tunnel-Interface, virtual-tunnel (vt) interfaces are created automatically to support ultimate-hop-popping upon enabling "protocol rsvp". These interfaces are associated with default IP and MPLS classifiers along with MPLS re-write rule. When "protocol rsvp" is disabled/enabled or MS-DPC/FPC (with tunnel-service) restarts, the vt interfaces are deleted and re-added to the system. However during the deletion,

these interfaces are not getting released from cosd process and thus leads to memory leak in cosd. [PR1071349](#)

### General Routing

- On MX104 router with SONET/SDH OC3/STM1 (Multi-Rate) MIC. In rare condition, if the MIC is plugged out from MX104, the Packet Forwarding Engine might crash, and the traffic forwarding will be affected. These MICs belong to SONET/SDH OC3/STM1 (Multi-Rate) MIC: \* MIC-3D-8OC3OC12-4OC48 \* MIC-3D-4OC3OC12-1OC48 \* MIC-3D-8CHOC3-4CHOC12 \* MIC-3D-4CHOC3-2CHOC12 \* MIC-3D-8DS3-E3 \* MIC-3D-8CHDS3-E3-B \* MIC-3D-1OC192-XFP. [PR997821](#)
- On MX Series platforms with MS-MPC/MS-MIC, if the "dump-on-flow-control" configuration statement is configured, traffic loss and the mspmand process crash might be observed when the MS-PIC comes up with traffic. [PR1037086](#)
- If default-address-selection configuration statement is configured on MX-VC, VC-heartbeat connection between member chassis may be unable to come up. [PR1041194](#)
- Queue stats on LSQ interfaces are not properly cleaned up when queuing is enabled on the IFD and the queues hosted at IFD level. This happens when there is a subsequent delete and create of LSQ interface (not always though). [PR1044340](#)
- On MX Series-based platform, when the feature flow-control is disabled (enabled by default) by using "no-flow-control" configuration statement (for example, under "gigether-options" hierarchy), after bringing up or rebooting the MPC, due to the fact that status of the hardware may not be updated correctly, the flow control on that MAC may remain enabled. [PR1045052](#)
- In subscriber management environment, the Berkeley Database (DB) may get into deadlock state. It is brought on by multiple daemons attempting to simultaneously access or update the same subscriber or service record. In this case, because the access to DB was blocked by device control daemon (dcd), the subscriber management infrastructure daemon (smid) fails to recover the DB. Consequently, the router may stop responding to all the login/logout requests as well as statistics activity. This timing-related issue is most likely to occur during login or logout and when the system is busy. [PR1054292](#)
- On MX Series routers, the interrupt-driven basis link down detection (an interrupt-driven link-down notification is generated to trigger locally attached systems to declare the interface down within a few milliseconds of failure) may fail after performing a unified in-service software upgrade (ISSU). The interrupt might have been prevented after performing unified ISSU because the interrupt registers were disabled before unified ISSU but never restored afterwards. [PR1059098](#)
- In an IPsec load-balancing environment using MS-MPC cards, the ICMP request and ICMP reply can go through two different IPsec tunnels due to asymmetric routing; that is, ICMP request goes through one PIC, and ICMP reply goes through another PIC. Because of this, the ICMP reply will get dropped and never reach the other side of the IPsec tunnel. [PR1059940](#)

- Due to incomplete fix, in releases containing PR869773 fix, rate limit drops are seen for Ingress queuing even though rate-limit is not configured or supported for ingress. [PR1061256](#)
- On MX Series router with MPC2E-3D-NG/MPC3E-3D-NG/MPC5/MPC6 linecards, the Ethernet frame loss measurement (ETH-LM) feature does not work. [PR1064994](#)
- When a route points to an aggregated multiservices (AMS) logical interface, then after manually bouncing this logical interface by disabling and then enabling it again, aggregate next hop referred by that route will have child unicast next hop pointing to .discard.0 interface instead of member interface (mams). As a result, traffic ingress on MPC card and routed to that route will be discarded. [PR1065944](#)
- If there are application-sets matching conditions in the NAT rule, NAT port might leak after deleting applications under application-set in live network. [PR1069642](#)
- With basic NAT44, when the router receiving packets on GRE tunnel, NAT was dropping all protocols other than PPTP on GRE tunnel. [PR1069872](#)
- Higher baseline CPU utilization and periodic CPU spikes might be seen on XM-based MPC as compared to MPC-3D-16XGE-SFPP cards due to the following reasons: On XM-based MPC, low priority threads which monitor various things in the background on a periodic basis such as voltage, temperature, stats counters, hardware status and so on are existed. When the system is idle, these threads are allowed to take more of the load, and that is why higher baseline CPU/CPU spikes are seen. This does not prevent other higher priority threads from running when they have to, as these are non-critical activities being done in the background and hence is a non impacting issue. [PR1071408](#)
- The overhead-accounting frame-mode command does not work on 100GbE CFP MIC, 100GbE CXP MIC, 2x40GbE QSFP MIC, and 10x10GbE SFPP MIC on MPC3E-3D-NG-Q, MPC3E-3D-NG, MPC2E-3D-NG-Q, and MPC2E-3D-NG. [PR1072001](#)
- This may be a false log message - the risk of false log is minor; however, the underlying error, for example, continuous fi recorder timeout, may impact traffic and can be major. When the specific log message is observed in the message file, please advise customer to investigate if there are continuous fabric errors, such as late cell, cell timeout and so on, on the reporting line card and recover those errors first. [PR1081771](#)
- MACsec using static secure association key (SAK) security mode does not work properly on MX80 routers and FPC slots other than slot 0 of MX104 routers. [PR1086117](#)
- On MX Series based line card, if a rlsq interface is receiving continuous fragmented traffic, doing rlsq switchovers couple of times might cause FPC to crash and reboot. [PR1088300](#)
- Some of the new revisions (for example, REV 30, REV 31) of the MICs cannot come up with NG-MPC2 or NG-MPC3 line card. We can check the MIC version by CLI command "show chassis hardware detail | no-more".  
root@user> show chassis hardware detail ..  
FPC 2 REV 14 750-054901 CADJ3871 MPC3E NG PQ & Flex Q CPU REV 11 711-045719  
CADN5465 RMPC PMB MIC 0 REV 30 750-028392 CAEB9203 3D 20x 1GE(LAN) SFP

Fan Tray 0 REV 05 740-014971 TP5127 Fan Tray Fan Tray 1 REV 05 740-014971 TP5103 Fan Tray. [PR1100073](#)

- Non-queuing MPC5E might crash continuously if rate-limit under transmit-rate for scheduler is applied. As a workaround, do not configure rate-limit and use firewall policer for forwarding-class instead. MPC5EQ is not exposed. [PR1104495](#)

### **Infrastructure**

- A reboot is needed if "chassis network services enhanced-ip" is configured on MX Series Universal Edge 3D Routers or on T4000 Routers with type 5 FPCs. Without the reboot, performing unified ISSU might cause the new master Routing Engine to crash and go to the db> prompt. [PR1013262](#)
- The issue was the gstatd for 64 bit was not getting to the correct path in the code and due to that gstat process was failing to start. [PR1074084](#)

### **Interfaces and Chassis**

- On dual Routing Engines platforms, as a High Availability (HA) method, master Routing Engine should relinquish mastership when both Routing Engine-to-Packet Forwarding Engine and Routing Engine-to-other-Routing Engine interfaces are down (this can be achieved only when GRES is enabled). But now on dual Routing Engine platforms except M10i and M20, master Routing Engine does not relinquish the mastership in such conditions, even executing CLI "request chassis routing-engine master acquire" on backup Routing Engine can not help. In such conditions, no FPC can be online without the connection to master Routing Engine. With the fix, the backup Routing Engine will take up the mastership automatically if both the internal link interfaces are down. [PR878227](#)
- On Ethernet PICs with longer hold down timer configured, flapping interface within the hold time might cause traffic loss longer than the hold period. [PR1040229](#)
- When configuring the Virtual Router Redundancy Protocol (VRRP) on an interface which is included in a routing-instance via applying groups setting, if changes are made to the interface, the VRRP process (vrrpd) memory leak might be observed on the device. [PR1049007](#)
- In Virtual Router Redundancy Protocol (VRRP) environment, after restarting the FPC, due to the Router Advertisement (RA) deletion is being incorrectly sent to routing protocol process (rpd) by VRRP process, the ICMPv6 may not be activated on the corresponding interfaces on the router that is acting as the master. In this case, no RA message could be sent out. [PR1051227](#)
- The "show chassis network-services" command might not show the correct configured value when executed on the backup Routing Engine. This command should only be executed on the master Routing Engine. [PR1054915](#)
- On DPC only chassis, after software upgrade or not graceful Routing Engine switchover, Ethernet OAM related LAG bundles might not come up due to the Link Fault Management (LFM) packets arrive on AE interface instead of physical link interface. [PR1054922](#)

- Two redundant logical tunnels (rlt) interfaces are configured with statement "per-unit-mac-disable" enabled. After configuring the second one, the first rlt interface goes down. `rlt0 { logical-tunnel-options { per-unit-mac-disable; <<<<<< } }` [PR1055005](#)
- The CLI description of the new 100-Gigabit Metro DWDM OTN PIC (PTX-2-100G-WDM-M) is different from the existing 100-Gigabit DWDM OTN PIC (P1-PTX-2-100G-WDM). The 100-Gigabit Metro DWDM OTN PIC's transceiver is identified as OTN-100G-M in the output from the show chassis hardware CLI command, and the cable type is identified as 100G METRO in the output from the show chassis pic CLI command. [PR1055325](#)
- There is a mismatch in mac statistics, few frames go unaccounted. This is a day-1 issue with the software fetching of mac statistics. The snap and clear bits were set together on pm3393 chip driver software, so it used to happen that even before the copy of stats to shadow registers happened, clear was happening which used to go unaccounted. Now rollover mechanism has been implemented and tested for 2 continuous days and everything is fine. [PR1056232](#)
- When "set chassis lcc 0 offline" is used on SCC and committed, the configuration gets synced on LCC. However, when "delete chassis lcc 0 offline" is used on SCC, we need to do commit two times on SCC in order to sync the configuration on LCC being brought online. [PR1058994](#)
- In multichassis link aggregation groups (MC-LAGs) environment, the MC-LAG peers have the MAC and port information and can forward the traffic appropriately. If a single VLAN on ICL interface is modified to a different VLAN, and then the administrator rolls back the VLAN configuration to the original one, the remote MAC might be stuck in the "Pending" state and not be installed in the bridge MAC-table, which causes the traffic forwarding to be affected. [PR1059453](#)
- When the Maximum Receive Unit (mru) value is not set under group-profile ppp-options hierarchy, a default value (1492) will be used. If mru value is set, the new value will take effect. But if the configured mru value is deleted from the group profile, the mru value remains the configured one and fails to fall back to the default one. [PR1059720](#)
- On MX Series routers, INET MTU (PPP payload MTU, that is IP header plus data excluding any L2 overhead) is being set to lowest MRU of either MX (local device) or peer. This behavior is not inline with ERX behavior, which is set to min(local MTU, peer MRU). This might cause the packet drops in the customer network in the downstream path. [PR1061155](#)
- In connectivity fault management (CFM) environment, if an AE interface is included in MEP interfaces, and if there is another AE interface configured without any child link (even this AE is not participating in OAM), the CFM sessions might not come up after Routing Engine restart or switchover. [PR1063962](#)
- Error message is continuously logged every second after a particular copper-SFP [P/N:740-013111] is plugged into a disabled port on MIC. \*\*\*\*\* error message \*\*\*\*\*  
mic\_sfp\_phy\_program\_phy: ge-\*/ \*/ - Fail to init PHY link mic\_periodic\_raw: MIC(\*/ \*)  
- Error in PHY periodic function PQ3\_IIC(WR): no target ack on byte 0 (wait spins 2)  
PQ3\_IIC(WR): I/O error (i2c\_stat=0xa3, i2c\_ctl[1]=0xb0, bus\_addr=0x56)  
mic\_i2c\_reg\_set - write fails with bus 86 reg 29 mic\_sfp\_phy\_write:MIC(\*/ \*) - Failed to write SFP PHY link 0, loc 29 mic\_sfp\_phy\_mdio\_sgmiilnk\_op: Failed to write: ifd = 140



ge-\*/\*/\*, phy\_addr: 0, phy\_reg: 29 ala88e1111\_reg\_write: Failed (20) to write register: phy\_addr 0x0, reg 0x1d Fails in function ala88e1111\_link\_init [PR1066951](#)

- To ensure that the router or switch is reachable for management purposes while it boots or if the routing protocol process fails to start properly, we can configure a backup router, which is a router that is directly connected to the local router or switch (that is, on the same subnet) through its private management interface (for example, fxp0 or me0). When a backup router running IPv6 and a static route to reach the management network are configured, some invalid IPv6 routes are added to default forwarding-table on the master or the backup Routing Engine. [PR1100981](#)

### **Layer 2 Features**

- BGP peer configured between two routers over lt (logical tunnel) interface, if deactivating and activating scaled configuration a few times, in rare condition, the lt interface might reject all the ARP reply packets, and hence the ARP resolution does not happen over this interface. Thus, the unicast routes are not in the correct states, and ping to such an lt interface will fail. [PR1059662](#)
- LACP partner system ID is shown incorrectly when the AE member link is connected to a different device, which might misguide while troubleshooting the LAG issues. [PR1075436](#)
- The Enhanced LAG feature is enabled in network-service enhanced-ip mode, but it is not supported in enhanced-ethernet mode. [PR1087982](#)

### **MPLS**

- The entropy label value allocated at times falls in the reserved mpls label range(0-15). The label value is calculated based on load balancing information and hence only certain mpls flows may encounter this issue. [PR1014263](#)
- With BGP labeled-unicast egress protection enabled in a Layer 3 VPN, the protected node advertises primary BGP labeled unicast routes that need protection. When there is next-hop change for a labeled route, for example, deactivating/activating egress-protection configuration statement or route churn, the memory might be exhausted which leads to the rpd process crash. [PR1061840](#)
- When fast-reroute, node-link-protection, or link-protection is configured, if a Shared Risk Link Group (SRLG) is associated with a link used by an LSP ingressing at a router, then on deleting the SRLG configuration from the router, the SRLG entry still stays in the SRLG table even after the re-optimization of this LSP. [PR1061988](#)
- When CSPF computes the path for node-protected bypass, it considers only the SRLG group configured on next-hop interface along the primary path. However it doesn't consider the SRLG group on next-to-next-hop interface to adequately provide diverse path between primary and node-protected bypass. [PR1068197](#)
- When a primary LSP gets re-routed due to better metric, Link/Node protection for this LSP is expected to come up within 7 seconds provided the bypass-lsp protecting the next-hop link/node is already available. However in some corner cases, the Link/Node protection for re-routed primary LSP will not come up within 7 seconds even with

bypass-lsp availability. The PR fixes this issue and reduces the delay of associating bypass-lsp with primary-lsp from 7 seconds to 2 seconds. [PR1072781](#)

- In MPLS environment, if one of minimum-signaling-bandwidth/merging-bandwidth/splitting-bandwidth/maximum-signaling-bandwidth is configured, or derived as value 0, the routing protocol process (rpd) may crash when lsp-splitting or lsp-merging (for example, when the traffic comes up/down) occurs. As a workaround, due to the logic of the configuration statement, none of the following configuration statements could be configured or derived as zero, -merging-bandwidth -minimum-signaling-bandwidth -splitting-bandwidth -maximum-signaling-bandwidth [PR1074472](#)

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

- SNMP queries for LAG MIB tables while LAG child interface is flapping may cause mib2d to grow in size and eventually crash with a core file. Mib2d will restart and recover by itself. [PR1062177](#)
- The text string of the SNMP object "system.sysDescr.0" does not include the Junos OS version of the device and displays the version of the FreeBSD kernel running on the Routing Engine instead. [PR1073232](#)

### ***Platform and Infrastructure***

- Recurring local memory (LMEM) data errors may cause lookup chip on MX Series with FPC wedge and eventually FPC crash. [PR1033660](#)
- If several aggregates are configured with shared-bandwidth-policer and those aggregates share the same Packet Forwarding Engine for child member links and one member links flaps, all traffic might get policed and dropped. The traffic dropped might not be on the bundle whose child member link flapped. [PR1035845](#)
- Due to a defect in the Junos OS software, when a telnet user experiences some undefined network disconnect, .perm and .env files under /var/run are left behind. This scenario happens only under certain unknown ungraceful network disconnects. When considerable number of .perm/.env files get accumulated under /var/run, issue is seen with telnet users, that they are not able to perform permitted operations on the router, post-login. [PR1047609](#)
- With VLAN manipulation configured for Ethernet Services, incorrect frame length might be used for egress policing on MX Series routers with MPCs/MICs. Currently, the frame length calculation is inconsistent for different traffic topology: 1. In case traffic crossed the fabric, the frame length prior to output VLAN manipulation is used; 2. In case of local traffic, the frame length prior to input VLAN manipulation is used. Actually the length after output VLAN manipulation should always be used. [PR1064496](#)
- When performing unified in-service software upgrade (ISSU) on MX Series routers with unsupported MICs (for example, "MIC-3D-8OC3OC12-4OC48") equipped, the MPC might crash during the field-replaceable unit (FRU) upgrade process. For example, unified ISSU is supported only by the MICs listed here on Junos OS Release 14.2: MIC-3D-20GE-SFP MIC-3D-2XGE-XFP MIC-3D-4XGE-XFP MIC-3D-40GE-TX MIC-3D-8OC3-2OC12-ATM MIC3-3D-2X40GE-QSFPP MIC3-3D-10XGE-SFPP MIC3-3D-1X100GE-CXP MIC3-3D-1X100GE-CFP. [PR1065731](#)

- Firewall filters which have a prefix-action can't be configured under [edit logical-system <name> firewall family inet] because the Packet Forwarding Engine won't be programmed for the filter. [PR1067482](#)
- If with about 1M routes on MX Series router, there might be more than 1 second (about 1.3s) packets dark window during unified ISSU. [PR1070217](#)
- VPLS filter applied under forwarding-options might drop VPLS frame unexpectedly when it is coming from an lt- interface. [PR1071340](#)
- If port-mirroring and VRRP over ae-irb is configured in a bridge-domain, enabling the Distributed Periodic Packet Management Process (ppmd) for VRRP in this BD might cause the VRRP to flap. [PR1071341](#)
- When inline-sampling is enabled, in race conditions, if packet gets corrupted and the corrupted packet length shows 0, this may cause "PPE\_x Errors thread timeout error" and eventually cause MPC card to crash. [PR1072136](#)
- VRRP advertisements might be dropped after enabling delegate-processing on the logical tunnel (lt) interface. It would result in VRRP master state observed on both routers. [PR1073090](#)
- When an MX Series chassis network-services is "enhanced-ip" and an AE with "family bridge" configuration is first committed, there is a possibility that an incorrect forwarding path may be installed causing traffic loss. [PR1081999](#)
- Issue is specific to 64-bit rpd and config-groups wildcard configuration specifically as in the following case: set groups TEST routing-instances <\*> routing-options multicast forwarding-cache family inet threshold suppress 200 set routing-instances vrf1 apply-groups TEST set routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 600. With this daemon(rpd) reads suppressed value "200" (that is, coming from groups) instead of reading value "600" from foreground, and customer sees unexpected behavior with respect to threshold-suppress. Workaround: They can replace wildcard with actual routing-instance name as in the following example: set groups TEST routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 200 set routing-instances vrf1 apply-groups TEST set routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 600 [PR1089994](#)

### ***Routing Protocols***

- Deletion of a routing-instance may lead to a routing daemon crash. This may happen if the routing-instance Routing Information Base (RIB) is referenced in an active policy-option configuration. As a workaround, when deactivating the routing-instance, all associated configurations using the route-table names in the routing-instance should also be deactivated. [PR1057431](#)
- In PIM environment, Bootstrap Router (BSR) can be used only between PIMv2 enabled devices. When deactivating all the interfaces which are running PIM bootstrap, the system changes to operate in PIMv1. At this time, all the information learned about/from the current BSR should be cleaned, but actually, BSR state is not cleaned. If the interface which was the previous "elected BSR" is activated, BSR state is PIM\_BSR\_ELECTED(should be cleaned previously) and the system assumes the BSR

timer is still here. When the system tries to access the null BSR timer, the rpd process might crash. [PR1062133](#)

- If with a large number of multicast sources for a same multicast group in PIM dense mode, the rpd process might crash after Routing Engine switchover. [PR1069805](#)
- For the pim nbr which is not directly connected (that is, nbr on unnumbered interface, or p2p interface with different subnet), pim join is not able to find the correct upstream nbr which results in join not propagating to the upstream nbr. show command for pim join shows upstream nbr "unknown". Issue is present in the 15.1R1 release. [PR1069896](#)
- In Protocol Independent Multicast (PIM) sparse mode environment, if the router is being used as the rendezvous point (RP) and also the last -hop router, when the (\*G) entry is present on the RP and a discard multicast route (for example, due to receiving multicast traffic from a non-RPF interface) is already existed, if the (S,G) entry is learned after receiving source-active (SA) of the Multicast Source Discovery Protocol (MSDP), the SPT cutover may fail to be triggered. There is no traffic impact as receivers still can get the traffic due to (\*G) route. [PR1073773](#)
- In multi-topologies IS-IS scenario, there is huge difference between estimated free bytes and actual free bytes when generating LSP with IPv6 prefix. It might cause LSP fragment exhaustion. [PR1074891](#)
- With Multicast Source Discovery Protocol (MSDP) and nonstop active routing (NSR) configured on the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP), the rpd process might permanently get stuck when multicast traffic received shortly after Routing Engines switchover. [PR1083385](#)
- 1. Configure the ospf and ospf3 in all routers 2. Configure node protection 3. Check for 22.1.1.0 any backup is present 4. Enable pplfa all 5. Check for 22.1.1.0 any pplfa backup is present through r2. We are not seeing any pplfa backup for 22.1.1.0. [PR1085029](#)

### **Services Applications**

- The session-limit-per-prefix feature for the MX Series DS-Lite server does not take Software flow into account when calculating the flow limit. [PR1023439](#)
- On MX Series routers with Multiservices 100, Multiservices 400, or Multiservices 500 PICs with "dump-on-flow-control" configured, if prolonged flow control failure, the coredump file might generate failure. [PR1039340](#)
- On MX Series routers that are acting as LNS to provide tunnel endpoints, it is observed that the service-interfaces are not usable if a MIC corresponding to them is not physically installed on the FPC. If only those service interfaces that belong to the removed PIC are added to service-device-pool, this results in no LNS subscribers being able to log in. Note that once the MIC is inserted into the FPC, the features could be used. [PR1063024](#)
- When configuring RADIUS authentication for Layer 2 Tunneling Protocol (L2TP), the RADIUS server cannot be recognized because the source address is not being read correctly. As a result, the L2TP session cannot be established. [PR1064817](#)
- The trigger for the crash is when the MS-DPCs Service PIC is in a low memory zone and it receives two SYN messages from the the same client IP within a very short time gap

in between the two SYNs. So this race condition is tied to running out of memory, failing to allocating a timer for a conversation, and having rapid SYNs on a TCP connection where the second TCP SYN is matched on flow which is being deleted due to a failed timer allocation for that. This scenario is very difficult to hit and should not be seen in production often. [PR1069006](#)

- Service PIC daemon (spd) might crash with core-dumps due to CGNAT pool's snmp-trap-thresholds configuration. [PR1070370](#)
- Earlier output from "show service l2tp tunnel" will not display tunnels with no sessions. This behavior have been changed, now empty tunnels are also displayed in this command. [PR1071923](#)

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

- Add "on <host>" argument to "request system software validate" to allow validation on a remote host/Routing Engine running Junos OS. [PR1066150](#)

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

- Due to a change in an existing PR, group names in the configuration must be a string of alphanumericals, dashes, or underscores. There is no workaround other than following the group name instructions. [PR1087051](#)

### ***VPNs***

- In the l2circuit environment, when l2ckt configuration has backup-neighbor, the flow-label operation is blocked at the configuration level. [PR1056777](#)
- On dual Routing Engines, if MVPN protocol itself is not configured, and nonstop active routing is enabled, the show command "show task replication" on the master Routing Engine will list the MVPN protocol even though it is not configured. Other than the misleading show output which may be slightly confusing to the user/customer, there is no functional impact due to this issue as such. There is no workaround available. [PR1078305](#)

- See Also**
- *New and Changed Features*
  - *Changes in Behavior and Syntax*
  - *Known Behavior*
  - *Known Issues*
  - *Documentation Updates*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Product Compatibility*

## Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1F7 documentation for the MX Series.

- [Subscriber Management Provisioning Guide on page 166](#)

---

### Subscriber Management Provisioning Guide

- The “enhanced-policer” topic erroneously states that when you commit a configuration that includes this statement, the CLI displays a warning that FPCs must be restarted for it to take effect, and prompts you to proceed with a restart. No such warning or prompt is displayed; instead, a warning message is logged that states that the enhanced policer is enabled on FPCs only after they are restarted.
- The topic “Configuring Address-Assignment Pool Linking” states that when you link multiple address-assignment pools, a secondary pool is used only when the primary address-assignment pool is fully allocated. However, once the router switches to a pool other than the primary, it continues using that pool even when addresses are available again in the primary pool.

- See Also**
- *New and Changed Features*
  - *Changes in Behavior and Syntax*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Known Issues*
  - *Resolved Issues*
  - *Migration, Upgrade, and Downgrade Instructions*
  - *Product Compatibility*

## Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the MX Series. Upgrading or downgrading Junos OS can take several minutes, depending on the size and configuration of the network.

Starting with Junos OS Release 15.1, in some of the devices, FreeBSD 10.x is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.1 remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see [Understanding Junos OS with Upgraded FreeBSD](#).



**NOTE:** Starting in Junos OS Release 15.1F4, Junos OS (FreeBSD 10.x) is also available to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to use the limited encryption Junos image (“Junos Limited”) for the FreeBSD 10.x Junos OS.

---



**CAUTION:** This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for MX Series routers. We do not recommend using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 15.1. The `request system software validate in-service-upgrade` command, which allows the detection of any compatibility issues before actually issuing the `request system software in-service-upgrade` command to initiate unified ISSU, is not supported in Junos OS Release 15.1 while upgrading from earlier Junos OS releases.

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.1-based Junos OS	FreeBSD 10.x-based Junos OS
MX5, MX10, MX40, MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

- [Basic Procedure for Upgrading to Release 15.1F7 on page 167](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) on page 169](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 170](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 172](#)
- [Upgrading a Router with Redundant Routing Engines on page 173](#)
- [Upgrading Using Unified ISSU on page 173](#)
- [Downgrading from Release 15.1 on page 174](#)

### Basic Procedure for Upgrading to Release 15.1F7

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 [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).



---

**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 Administration Library](#).

---



### Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)

Products impacted: MX240, MX480, MX960, MX2010, and MX2020.

To download and install from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x):

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

<https://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 all the countries except Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia, use the following command:

For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-15.1F7.11.tgz
```

For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-15.1F7.11.tgz
```

- Customers in Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia, use the following command:

For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-32-15.1F7.11-limited.tgz
```

For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-15.1F7.11-limited.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)

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 10.x, and Junos OS (FreeBSD 6.1) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

Use the **reboot** command to reboot 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 15.1F7 **jinstall** package, you cannot issue the **request system software rollback** command to return to the previously installed Junos OS (FreeBSD 6.1) software. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

---

### Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)

Products impacted: MX5, MX10, MX40, MX80 and MX104.



**NOTE:** Customers in Belarus, Kazakhstan, and Russia must use the following procedure for all MX Series routers running Junos OS Release 15.1.

To download and install from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1):

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

<https://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-15.1F7.11-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-15.1F7.11-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.

Use the **reboot** command to reboot 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 15.1F4 jinstall package, you cannot issue the **request system software rollback** command to return to the previously installed software. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

---

### Upgrade 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 13.3, 14.1, and 14.2 are EEOL releases. You can upgrade from Junos OS Release 13.3 to Release 14.2 or even from Junos OS Release 13.3 to Release 15.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

---

### Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform the following Junos OS installation on each Routing Engine separately to avoid disrupting network operation:

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 [Installation and Upgrade Guide](#).

---

### Upgrading Using Unified ISSU



**CAUTION:** This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for MX Series routers. We do not recommend using unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 15.1.

---

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 ISSU 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 ISSU, see the [High Availability Feature Guide for Routing Devices](#).

#### ***Changes Planned for Future Releases***

Starting in Junos OS Release 15.1F6-S1 and the subsequent 15.1F6-S(x) releases (for example, 15.1F6-S2), ISSU would be supported on the following Modular Port Concentrators (MPCs):

- MX-MPC3E-3D
- MPC4E-3D-32XGE-SFPP
- MPC4E-3D-2CGE-8XGE
- MPC5E-100G10G

## Downgrading from Release 15.1

---

To downgrade from Release 15.1 to another supported release, follow the procedure for upgrading, but replace the 15.1 `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 *Installation and Upgrade Guide*.

- See Also**
- *New and Changed Features*
  - *Changes in Behavior and Syntax*
  - *Known Behavior*
  - *Known Issues*
  - *Resolved Issues*
  - *Documentation Updates*
  - *Product Compatibility*

## Product Compatibility

- [Hardware Compatibility on page 174](#)

### Hardware Compatibility

---

To obtain information about the components that are supported on the devices, and special compatibility guidelines with the release, see the Hardware Guide and the Interface Module Reference for the product.

To determine the features supported on MX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

<https://pathfinder.juniper.net/feature-explorer/>

- See Also**
- *New and Changed Features*
  - *Changes in Behavior and Syntax*
  - *Known Behavior*
  - *Known Issues*
  - *Resolved Issues*

- *Documentation Updates*
- *Migration, Upgrade, and Downgrade Instructions*

## Third-Party Components

---

This product includes third-party components. To obtain a complete list of third-party components, see [Overview for Routing Devices](#).

For a list of open source attributes for this Junos OS release, see [Open Source: Source Files and Attributions](#).

## Finding More Information

---

For the latest, most complete information about known and resolved issues with Junos OS, see the Juniper Networks Problem Report Search application at: <https://prsearch.juniper.net>.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>.

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF format. Find Content Explorer at: <https://www.juniper.net/documentation/content-applications/content-explorer/>.

## 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 post sales technical support, you can access our tools and resources online or open a case with JTAC.

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or Partner Support Service support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://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: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://www.juniper.net/kb/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

## 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 <https://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 <https://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/documentation/feedback/>.

---

## Revision History

27 February 2018—Revision 2, Junos OS Release 15.1F7— MX Series.

21 March 2017—Revision 1, Junos OS Release 15.1F7— MX Series.

Copyright © 2018 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be 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.