

Release Notes: Junos[®] OS Release 15.1F6 for the MX Series, PTX Series, and T Series

15.1F6
23 March 2017

Contents

Introduction	4
Junos OS Release Notes for MX Series 3D Universal Edge Routers and T Series	
Core Routers	4
New and Changed Features	5
Hardware	5
Class of Service (CoS)	9
General Routing	10
High Availability and Resiliency	12
Interfaces and Chassis	12
IPv6	14
Management	14
MPLS	14
Multicast	15
Network Management and Monitoring	15
Routing Protocols	16
Services Applications	19
Software Installation and Upgrade	21
Software-Defined Networking	21
Subscriber Management and Services	24
System Logging	31
VPNs	31
Changes in Behavior and Syntax	32
General Routing	33
Interfaces and Chassis	33
IPv6	34
MPLS	34
Network Management and Monitoring	34
Routing Policy and Firewall Filters	35
Routing Protocols	35
Services Applications	35
Subscriber Management and Services (MX Series)	35

System Logging	37
System Management	43
Virtual Chassis	43
VPNs	43
Known Behavior	44
Hardware	44
General Routing	44
Interfaces and Chassis	47
MPLS	47
OpenFlow	47
Subscriber Management and Services (MX Series)	47
Known Issues	48
Forwarding and Sampling	48
General Routing	48
High Availability and Resiliency	49
MPLS	49
Platform and Infrastructure	49
Routing Protocols	50
Services Applications	51
VPNs	51
Resolved Issues	51
Resolved Issues: 15.1F6	51
Resolved Issues: 15.1F5	86
Resolved Issues: 15.1F4	102
Resolved Issues: 15.1F3	107
Resolved Issues: 15.1F2	116
Documentation Updates	126
Subscriber Management Provisioning Guide	126
Migration, Upgrade, and Downgrade Instructions	126
Basic Procedure for Upgrading to Release 15.1F5	127
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1) ..	129
Upgrade and Downgrade Support Policy for Junos OS Releases	130
Upgrading a Router with Redundant Routing Engines	131
Upgrading the Software for a Routing Matrix	131
Upgrading Using Unified ISSU	132
Downgrading from Release 15.1	133
Product Compatibility	133
Hardware Compatibility	133
Junos OS Release Notes for PTX Series Packet Transport Routers	135
New and Changed Features	135
Class of Service	136
Hardware	136
General Routing	141
Interfaces and Chassis	141
Management	143
MPLS	143
Multicast	144
Network Management and Monitoring	144
Routing Policy and Firewall Filters	144

Routing Protocols	145
Services Applications	146
Software-Defined Networking	147
User Interface and Configuration	148
VPNs	148
Changes in Behavior and Syntax	149
Hardware	150
Forwarding and Sampling	150
General Routing	150
IPv6	150
Network Management and Monitoring	150
Routing Policy and Firewall Filters	151
Known Behavior	151
General Routing	151
Known Issues	154
General Routing	154
Integrated Photonic Line Card and Expansion Module	154
MPLS	155
Routing Protocols	155
Resolved Issues	156
Resolved Issues: 15.1F6	156
Resolved Issues: 15.1F5	161
Resolved Issues: 15.1F4	163
Resolved Issues: 15.1F3	165
Resolved Issues: 15.1F2	167
Migration, Upgrade, and Downgrade Instructions	169
Upgrading Using Unified ISSU	169
Upgrading a Router with Redundant Routing Engines	169
Basic Procedure for Upgrading to Release 15.1F5	170
Product Compatibility	172
Hardware Compatibility	173
Third-Party Components	174
Finding More Information	174
Documentation Feedback	174
Requesting Technical Support	175
Self-Help Online Tools and Resources	175
Opening a Case with JTAC	175
Revision History	176

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.

These release notes accompany Junos OS Release 15.1F6 for the MX Series, PTX Series and T 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 and T Series Core Routers

These release notes accompany Junos OS Release 15.1F6 for the MX Series and T 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 <http://www.juniper.net/techpubs/software/junos/>.



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



NOTE: The following T Series routers are not supported in 15.1F6:

- T320
- T640
- T1600

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 32](#)
- [Known Behavior on page 44](#)
- [Known Issues on page 48](#)
- [Resolved Issues on page 51](#)
- [Documentation Updates on page 126](#)
- [Migration, Upgrade, and Downgrade Instructions on page 126](#)
- [Product Compatibility on page 133](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F6 for the MX Series and the T Series.

- [Hardware on page 5](#)
- [Class of Service \(CoS\) on page 9](#)
- [General Routing on page 10](#)
- [High Availability and Resiliency on page 12](#)
- [Interfaces and Chassis on page 12](#)
- [IPv6 on page 14](#)
- [Management on page 14](#)
- [MPLS on page 14](#)
- [Multicast on page 15](#)
- [Network Management and Monitoring on page 15](#)
- [Routing Protocols on page 16](#)
- [Services Applications on page 19](#)
- [Software Installation and Upgrade on page 21](#)
- [Software-Defined Networking on page 21](#)
- [Subscriber Management and Services on page 24](#)
- [System Logging on page 31](#)
- [VPNs on page 31](#)

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.

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.

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 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.

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, jdhcpd, 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
```

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

**Related
Documentation**

- *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 15.1F6 for the MX Series and T Series.

- [General Routing on page 33](#)
- [Interfaces and Chassis on page 33](#)
- [IPv6 on page 34](#)
- [MPLS on page 34](#)
- [Network Management and Monitoring on page 34](#)
- [Routing Policy and Firewall Filters on page 35](#)
- [Routing Protocols on page 35](#)
- [Services Applications on page 35](#)
- [Subscriber Management and Services \(MX Series\) on page 35](#)
- [System Logging on page 37](#)

- [System Management on page 43](#)
- [Virtual Chassis on page 43](#)
- [VPNs on page 43](#)

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**.

Interfaces and Chassis

- **Support for fabric self-pings and Packet Forwarding Engine liveness in single-chassis systems (T Series)**—In Junos OS Release 15.1 F6, T Series single-chassis systems support the fabric self-ping and Packet Forwarding Engine liveness mechanisms to detect fabric degradation and avoid a traffic black hole. If any error is detected by these two mechanisms, the fabric manager raises a *fabric degraded alarm* and initiates recovery by restarting the FPC. In a single-chassis system, FPC restart is enabled by default, unlike in a multichassis system where FPC restart is disabled by default.
- **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](#).]

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.

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.

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.

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	<i>software-string src-ip:src-port [xlated-src-ip:xlated-src-port]->[xlated-dst-ip: xlated-dst-port]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	<i>software-string src-ip:src-port [xlated-src-ip:xlated-src-port]->[xlated-dst-ip: xlated-dst-port]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	<i>software-string src-ip:src-port [xlated-src-ip:xlated-src-port]->[xlated-dst-ip: xlated-dst-port]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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The ICMP error packet was dropped because it did not belong to an existing flow.	LOG_NOTICE
JSERVICES_ICMP_HEADER_LEN_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the contents of two fragments overlapped.	LOG_NOTICE
JSERVICES_IP_PACKET_CHECKSUM_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because checksum was incorrect.	LOG_NOTICE
JSERVICES_IP_PACKET_DST_BAD	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the length of a fragment was invalid.	LOG_NOTICE
JSERVICES_IP_PACKET_INCORRECT_LEN	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The IP packet is discarded because packet length was invalid.	LOG_NOTICE
JSERVICES_IP_PACKET_LAND_ATTACK	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet version was not IPv4.	LOG_NOTICE
JSERVICES_IP_PACKET_NOT_VERSION_6	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet version was not IPv6.	LOG_NOTICE
JSERVICES_IP_PACKET_PROTOCOL_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because it used invalid IP protocol.	LOG_NOTICE
JSERVICES_IP_PACKET_SRC_BAD	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet did not contain the minimum amount of data required.	LOG_NOTICE
JSERVICES_NO_IP_PACKET	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	Packet received was not an IPv4 or IPv6 packet.	LOG_NOTICE
JSERVICES_SYN_DEFENSE	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	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	<i>source-ip:destination-ip</i> 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 -> 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 -> 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 -> 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 -> 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 -> 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 -> 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 -> 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 -> 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 -> 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.

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
```

Related Documentation

- [New and Changed Features on page 5](#)
- [Migration, Upgrade, and Downgrade Instructions on page 126](#)
- [Known Issues on page 48](#)
- [Resolved Issues on page 51](#)
- [Documentation Updates on page 126](#)
- [Migration, Upgrade, and Downgrade Instructions on page 126](#)
- [Product Compatibility on page 133](#)

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1F6 for the MX Series and T Series.

- [Hardware on page 44](#)
- [General Routing on page 44](#)
- [Interfaces and Chassis on page 47](#)
- [MPLS on page 47](#)
- [OpenFlow on page 47](#)
- [Subscriber Management and Services \(MX Series\) on page 47](#)

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.

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.

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.

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
```

**Related
Documentation**

- *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 15.1F6 for the MX Series and T Series.

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

- [Forwarding and Sampling on page 48](#)
- [General Routing on page 48](#)
- [High Availability and Resiliency on page 49](#)
- [MPLS on page 49](#)
- [Platform and Infrastructure on page 49](#)
- [Routing Protocols on page 50](#)
- [Services Applications on page 51](#)
- [VPNs on page 51](#)

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 policer even if the traffic is not exceeding the configured bandwidth. As a workaround, configure the "shared-bandwidth-policer" configuration statement under the policer. [PR1125071](#)

General Routing

- The following message is generated every 5 second in MX104 on 14.2R1~R3 and 15.1R1: "xxx chassisd[1362]: Cannot read hw.chassis.startup_time: No such file or directory." [PR1049015](#)
- Certain VTY JNH commands(see description of this PR-1094955) on Trinity platform will not decode properly, would need this PR fix. [PR1094955](#)
- SFB2 offline/online with 20 line cards takes 9 minutes 52 seconds whereas for SFB it takes 42 seconds. [PR1097338](#)

- On MX Series platform, executing CLI command "request ancp oam neighbor" might cause the ancpd process crash. [PR1125230](#)
- Subscriber where TCP is attached to the underlying IFL will errantly end up in the control IFL queue. Workaround is to attach a TCP profile to each subscriber IFL. [PR1162108](#)
- Under loaded router (i.e., 1000+ IPsec tunnels, 100+ BGP sessions, 1+ Mln routes) when performing a GRES some IPsec tunnels might fail to come up. [PR1162385](#)
- Default EVPN policy in junos-defaults for mx-series is removed. This was used to enable per packet load-balance for EVPN routes. Now per-packet load balance needs to be configured explicitly. [PR1162433](#)
- Stacked ifl and the underlying ifl cannot be part of the same iflset. [PR1162805](#)
- On MS-MIC, starting from 15.1R3 onwards, the J-Flow scaling has come down to 12.5 M active flows. [PR1163976](#)
- Traffic may drop during Routing Engine switchover. [PR1164107](#)

High Availability and Resiliency

- After graceful switchover is triggered in the master VRRP router for the first time, the master state for all the VRRP instances is toggled to backup and comes back to master immediately. During this time all the traffic is dropped and comes back. [PR1142227](#)
- 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](#)

MPLS

- These benign error messages can be ignored. The Junos OS code will be cleaned up to remove these message in later code. [PR1136033](#)
- FPC cores may be seen during AE flap on Type 5-3D FPCs. [PR1164175](#)

Platform and Infrastructure

- Once the Traffic Offload Engine (TOE) thread is stalled due to memory error at the lookup chip all statistics collection from the interfaces hosted by this PFE are not updated anymore. [PR1051076](#)
- On MX Series Virtual Chassis (MX-VC) with "locality-bias" configured, when equal-cost multipath (ECMP) load-balancing is occurring in the VC system, multicast streams and flooded Layer 2 streams may be duplicated or lost. As a workaround, we can disable "locality-bias" if possible. [PR1104096](#)
- When replying to operational RPCs executed over NETCONF sessions, the Junos OS software will add leading and trailing whitespace that is not present for Junoscript sessions. [PR1143761](#)
- This issue can occur in subscriber management configuration that has a very large number of address pools on the order of 64K or more. At this scale of address pools, synchronizing the two Routing Engines after one GRES switchover can take a long time

preventing a second graceful switchover. This issue is not expected in configurations with smaller address pool scale. [PR1159972](#)

- VMX PE could send ICMP packet source from address of a disabled CE interface in L3 VRF instance. This could break the ping test without specifying the source address. An example is shown below that VMX PE send ICMP packet sourced from ae0.306 (disabled): 202.202.6.1. ae0.301 up up inet 202.202.1.1/30 ae0.306 down down inet 202.202.6.1/30 Router> ping 14.1.1.2 routing-instance VPN1 rapid PING 14.1.1.2 (14.1.1.2): 56 data bytes --- 14.1.1.2 ping statistics --- 5 packets transmitted, 0 packets received, 100% packet loss 15:54:57.141675 In IP 202.202.6.1 > 14.1.1.2: ICMP echo request, id 14135, seq 0, length 64 15:54:57.141697 Out IP 14.1.1.2 > 202.202.6.1: ICMP echo reply, id 14135, seq 0, length 64 Workaround is to enable the ae0.306, or just add a lo0 logical interface under this L3 VRF instance VPN1, which can be used as ICMP source by default. [PR1175658](#)

Routing Protocols

- The static/static access routes pointing to an unnumbered interface are getting added in the routing table even if the interface is down. In this case, if graceful Routing Engine switchover (GRES) is disabled, this type of route will never be added in routing table after Routing Engine switchover. [PR1064331](#)
- When BGP speaker has multiple peers configured in a BGP group and when it receives the route from a peer and re-advertises route to another peer within the same group, MIB object "jnxBgpM2PrefixOutPrefixes" to the peers in the same group reports the total number of advertised prefixes in the group. MIB value "jnxBgpM2PrefixOutPrefixes" is defined as per peer basis but it looks as if it is per group basis. As a workaround, we can get the number of advertised prefixes from CLI command "show bgp neighbor" instead. [PR1116382](#)
- When multiple addresses are configured on an interface, if the interface has "interface-type p2p" configured under OSPF and the router does not receive any OSPF packets from one of the IFAs, the OSPF state will not go down for the corresponding adjacency. It should have no impact on route learning, but it might cause confusion for troubleshooting, when peering with Cisco devices, which have multiple addresses configured as secondary addresses. [PR1119685](#)
- On dual-Routing Engine platforms, in scaling scenario (e.g., there are 6 million routes on "old" master Routing Engine), if graceful Routing Engine switchover (GRES) or graceful restart (GR) is not enabled, the routing protocol process (rpd) may crash on the "old" master Routing Engine after performing Routing Engine switchover. As a workaround, if possible, rebooting the "old" master Routing Engine (new backup Routing Engine) after switchover could avoid the issue. [PR1128023](#)
- Generate route does not inherit the next-hop from the contributing route in L3VPN case when the contributing route is learnt via MP-BGP. The next-hop remains as reject for the generated route. [PR1149970](#)

Services Applications

- Space may be missing in tnp.bootpd log message output string. There is no known operational impact. [PR1075355](#)

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 (mvprinstancename.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](#)

Related Documentation

- *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 15.1F6 for the MX Series and T Series.

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

- [Resolved Issues: 15.1F6 on page 51](#)
- [Resolved Issues: 15.1F5 on page 86](#)
- [Resolved Issues: 15.1F4 on page 102](#)
- [Resolved Issues: 15.1F3 on page 107](#)
- [Resolved Issues: 15.1F2 on page 116](#)

Resolved Issues: 15.1F6

- [Class of Service \(CoS\) on page 52](#)
- [Forwarding and Sampling on page 52](#)
- [General Routing on page 54](#)
- [High Availability \(HA\) and Resiliency on page 67](#)
- [Infrastructure on page 67](#)

- [Interfaces and Chassis on page 68](#)
- [Layer 2 Features on page 70](#)
- [MPLS on page 71](#)
- [Network Management and Monitoring on page 73](#)
- [Platform and Infrastructure on page 74](#)
- [Routing Policy and Firewall Filters on page 80](#)
- [Routing Protocols on page 80](#)
- [Services Applications on page 84](#)
- [Subscriber Management and Services on page 85](#)
- [User Interface and Configuration on page 85](#)
- [VPNs on page 86](#)

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 platform 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: rtslib: 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 rtslib will not be received at pfed. (Please note, rtslib broadcasts events related to interface add, delete, modified etc. The pfe clients registered itself with rtslib library to listen for these events. In this particular scenario, rtslib 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 0 100.100.100.3 <<<<< bfd packet 299776 13 0 0.0.0.17
0.1.134.160 0 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 LIBTASK memory subsystem. This PR impacts all daemons using LIBTASK (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 msec 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 T Series multichassis platforms, when offline and then online the LCC from SCC (e.g., executing the CLI command "set chassis lcc 0 offline" command, and then executing "delete set chassis lcc 0 offline") in quick successions (that is, within the timeout setting for peer to reconnect, 60 seconds, which is not configurable), kernel replication error "ENOENT" may occur, which can cause ksyncd to crash and in thus trigger a live vmcore. Additionally, this is a timing issue and LCC offline followed by online within 60 seconds is the only known trigger so far. As a workaround, on the safer side, it is recommended to online the LCC after 120 seconds. [PR1108048](#)
- 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 platform 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 platform, 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](#)
- User now has access to common DHCP CLI commands on VMX, including - Show dhcp client binding - Show dhcp client statistics - Clear dhcp client binding - Clear dhcp client statistics - Request dhcp client renew (To renew lease). [PR1152115](#)
- 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. [http://www.juniper.net/documentation/en_US/junos15.1/topics/reference/state ment-hierarchy/system-services-resource-monitor.html](http://www.juniper.net/documentation/en_US/junos15.1/topics/reference/state%20ment-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](#)
- "op 8 (COS Blob) failed" messages may be seen in syslog for vmx when we reboot the FPC. [PR1156450](#)
- 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](#)

- In the TXP environment, the Line-Card Chassis (LCC) Switch Interface Board (SIB) status is not right when execute command "user@router> show chassis environment", their status are Absent, but no alarms. This is a minor issue, it does not affect business. [PR1156841](#)
- 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 platform 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 15.2IB, if deactive 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](#)
- In T-series platform in 15.1Fx, the error log message about "IFD Ether boolean set (opcode 55) failed" and "SXGE: unknown option 132" might be seen on certain scenarios such as: RE restarting, FPC restarting or RE switchover. [PR1173707](#)
- 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](#)
- Starting with 15.1F5, the splitting of destination NAT pools across AMS members will be prevented. Currently with AMS interfaces, dnat44 pools do not get split. However, all twice-NAT destination pools are split. This is not needed and this change makes it so (source pools are split or/and hashing is based on source so there is never any chance of conflict). [PR1184749](#)

- [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:
<http://kb.juniper.net/InfoCenter/KB9382>. [PR1144093](#)

- In MX-VC or VRR 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 platform 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 platform, 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 platform 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 be 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](#)
- When polling SNMP IF-MIB table on VMX platforms, the unicast packet counters, such as "ifInUcastPkts" and "ifHCOOutUcastPkts" are always zero for IFD (port- level) interfaces. [PR1155895](#)
- 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 platform, when offlining 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-IOC48
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 PFE 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 platform 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 RE 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 1HASH_OP not done 0x0300a87601beff03 0000000000000001. 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/Router Engine CPU. [PR110854](#)
- 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 RE core. The exact root cause is not known. While processing updates on the backup RE (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 <http://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 sequence (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 MVPN 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 87](#)
- [Class of Service \(CoS\) on page 92](#)
- [Forwarding and Sampling on page 92](#)

- [High Availability \(HA\) and Resiliency on page 92](#)
- [Infrastructure on page 93](#)
- [Interfaces and Chassis on page 93](#)
- [Layer 2 Features on page 96](#)
- [MPLS on page 96](#)
- [Network Management and Monitoring on page 97](#)
- [Platform and Infrastructure on page 97](#)
- [Routing Policy and Firewall Filters on page 99](#)
- [Routing Protocols on page 100](#)
- [User Interface and Configuration on page 101](#)
- [VPNs on page 101](#)

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 ($64 \times 8 = 512$ 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 mspppmond 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 M Series, MX Series, T 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/TFTP/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](#)
- On VMX platforms with Junos OS Release 14.1R5 and later, when traffic runs over an extended period of time (e.g., 10 hours), Scheduler Oinker messages might be seen on fpc 0 for various threads. The threads are mostly Packet Forwarding Engine Manager, COS HALP, and Idle. This is unlikely to affect traffic. [PR1140360](#)

- 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](#)
- On MX Series platforms with FPC3, the octets of IPv4 source and destination addresses in the firewall log are listed reverse; this might affect troubleshooting. The IPv6 log works fine. This is a minor issue, there is no other service impact. [PR1141495](#)
- 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 an 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](#)
- "op 8 (COS Blob) failed" messages might be seen in the syslog for vmx when rebooting the FPC. [PR1156450](#)

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 fmgAlarmTime. 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. [PR1116826](#)
- 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 verge. [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:
<http://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 vrpd 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-10C192-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](#)
- When polling SNMP IF-MIB table on VMX platforms, the unicast packet counters, such as "ifInUcastPkts" and "ifHCOutUcastPkts" are always zero for IFD (port- level) interfaces. [PR1155895](#)

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 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](#)
- 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 102](#)
- [Forwarding and Sampling on page 102](#)
- [General Routing on page 102](#)
- [Interfaces and Chassis on page 104](#)
- [Layer 2 Features on page 104](#)
- [MPLS on page 105](#)
- [Network Management and Monitoring on page 105](#)
- [Platform and Infrastructure on page 105](#)
- [Routing Protocols on page 106](#)
- [Services Applications on page 107](#)
- [VPNs on page 107](#)

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. [PR110590](#)
- 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](#)
- This is a cosmetic issue that vMX firewall logs may show wrong packet length for dropped packets. [PR1124855](#)
- 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 0call_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 107](#)
- [Infrastructure on page 111](#)
- [Interfaces and Chassis on page 112](#)
- [Layer 2 Features on page 112](#)
- [MPLS on page 113](#)
- [Network Management and Monitoring on page 113](#)
- [Platform and Infrastructure on page 113](#)
- [Routing Protocols on page 116](#)
- [Software Installation and Upgrade on page 116](#)
- [VPNs on page 116](#)

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](#)
- When VMX is deployed, initially there is no management port configuration, so configuration needs to be applied by serial console. The console for VMX is set to 9600

baud rate. With this rate, only a small number of configuration lines can be pasted at a time. [PR1068152](#)

- 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](#)
- High latency might be observed when continuous IPv6 pings are sent to VMX platform. [PR1096403](#)
- 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 platform 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 platform 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/ift) 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 (ift) 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> [PR1112580](#)

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:

http://kb.juniper.net/InfoCenter/index?page=content&id=JSA10701&cat=SIRT_1&actp=LIST. [PR1102581](#)

- 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](#)
- On T4000 platform with FPC Type-5 equipped, after performing unified ISSU, due to the fact that only 6 out of 16 temperature sensors may get initialized, the temperature reading for the line card may be shown as "Absent". [PR1104240](#)
- 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 117](#)
- [General Routing on page 117](#)
- [Infrastructure on page 119](#)
- [Interfaces and Chassis on page 119](#)
- [Layer 2 Features on page 121](#)
- [MPLS on page 121](#)
- [Network Management and Monitoring on page 122](#)
- [Platform and Infrastructure on page 122](#)
- [Routing Protocols on page 124](#)

- [Services Applications on page 125](#)
- [Software Installation and Upgrade on page 125](#)
- [User Interface and Configuration on page 125](#)
- [VPNs on page 125](#)

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 platform 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 platform 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 | no-more
Hardware inventory:
Item Version Part 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
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_sgmii_lnk_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](#)

- For a Routing Matrix, if different Routing Engine models are used on switch-card chassis (SCC)/switch-fabric chassis (SFC) and line-card chassis (LCC) (for example, RE-1600 on SCC/SFC and RE-DUO-C1800 on LCC), where the out-of-band (OoB) management interfaces are named differently (for example, fxp0 on SCC/SFC Routing Engine and em0 on LCC Routing Engine), then the OoB management interface configuration for LCC Routing Engine will not be propagated from SCC/SFC Routing Engine during commit. [PR1050743](#)
- 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 Softwire flow into account when calculating the flow limit. [PR1023439](#)
- On MX Series routers and T 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](#)

- Related Documentation**
- *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 15.1F6 documentation for the MX Series, and T Series.

- [Subscriber Management Provisioning Guide on page 126](#)

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.

- Related Documentation**
- [New and Changed Features on page 5](#)
 - [Changes in Behavior and Syntax on page 32](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 126](#)
 - [Known Issues on page 48](#)
 - [Resolved Issues on page 51](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 126](#)
 - [Product Compatibility on page 133](#)

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for and T 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 T 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.

- [Basic Procedure for Upgrading to Release 15.1F5 on page 127](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 129](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 130](#)
- [Upgrading a Router with Redundant Routing Engines on page 131](#)
- [Upgrading the Software for a Routing Matrix on page 131](#)
- [Upgrading Using Unified ISSU on page 132](#)
- [Downgrading from Release 15.1 on page 133](#)

Basic Procedure for Upgrading to Release 15.1F5

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 for Routing Devices](#).

.....

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

Products impacted: All T Series routers and the 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:
<http://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.



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

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

```
user@host> request system software add validate reboot
source/jinstall-15.1F5.11-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-15.1F5.11-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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform 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 the Software for a Routing Matrix

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

- A minimum of free disk space and DRAM on each Routing Engine. The software upgrade will fail on any Routing Engine without the required amount of free disk space and DRAM. To determine the amount of disk space currently available on all Routing Engines of the routing matrix, use the CLI **show system storage** command. To determine the amount of DRAM currently available on all the Routing Engines in the routing matrix, use the CLI **show chassis routing-engine** command.
- The master Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC are all re0 or are all re1.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC are all re1 or are all re0.
- All master Routing Engines in all routers run the same version of software. This is necessary for the routing matrix to operate.
- All master and backup Routing Engines run the same version of software before beginning the upgrade procedure. Different versions of Junos OS can have incompatible message formats especially if you turn on GRES. Because the steps in the process include changing mastership, running the same version of software is recommended.
- For a routing matrix with a TX Matrix router, the same Routing Engine model is used within a TX Matrix router (SCC) and within a T640 router (LCC) of a routing matrix. For example, a routing matrix with an SCC using two RE-A-2000s and an LCC using two RE-1600s is supported. However, an SCC or an LCC with two different Routing

Engine models is not supported. We suggest that all Routing Engines be the same model throughout all routers in the routing matrix. To determine the Routing Engine type, use the CLI **show chassis hardware | match routing** command.

- For a routing matrix with a TX Matrix Plus router, the SFC contains two model RE-DUO-C2600-16G Routing Engines, and each LCC contains two model RE-DUO-C1800-8G or RE-DUO-C1800-16G Routing Engines.



BEST PRACTICE: Make sure that all master Routing Engines are re0 and all backup Routing Engines are re1 (or vice versa). For the purposes of this document, the master Routing Engine is re0 and the backup Routing Engine is re1.

To upgrade the software for a routing matrix:

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

For the detailed procedure, see the [Routing Matrix with a TX Matrix Router Deployment Guide](#) or the [Routing Matrix with a TX Matrix Plus Router Deployment 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 and T 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 **install** package with one that corresponds to the appropriate release.



NOTE: You cannot downgrade more than three releases. For example, if your routing platform is running Junos 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*.

Related Documentation

- *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 133](#)

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:

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

**Related
Documentation**

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

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 15.1F6 for the PTX 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 <http://www.juniper.net/techpubs/software/junos/>.



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

- [New and Changed Features on page 135](#)
- [Changes in Behavior and Syntax on page 149](#)
- [Known Behavior on page 151](#)
- [Known Issues on page 154](#)
- [Resolved Issues on page 156](#)
- [Migration, Upgrade, and Downgrade Instructions on page 169](#)
- [Product Compatibility on page 172](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F6 for the PTX Series.

- [Class of Service on page 136](#)
- [Hardware on page 136](#)
- [General Routing on page 141](#)
- [Interfaces and Chassis on page 141](#)
- [Management on page 143](#)
- [MPLS on page 143](#)
- [Multicast on page 144](#)
- [Network Management and Monitoring on page 144](#)
- [Routing Policy and Firewall Filters on page 144](#)
- [Routing Protocols on page 145](#)
- [Services Applications on page 146](#)
- [Software-Defined Networking on page 147](#)
- [User Interface and Configuration on page 148](#)
- [VPNs on page 148](#)

Class of Service

- **Change in scaling number for rewrite rules (PTX Series)**—Starting with Release 15.1F2, on PTX Series routers, the scaling number for a rewrite rule is reduced by one when the default EXP rewrite is used. This change in scaling number is introduced to:
 - Support all possible combinations of EXP rewrite rules.
 - Fix the issue of incorrect modification of EXP bits of the inner label by the default MPLS EXP rewrite rule during the label pop operation.

Hardware

- **New Routing Engine RE-PTX-X8-64G (PTX5000)**—Starting in Junos OS Release 15.1F4, the Routing Engine RE-PTX-X8-64G is supported on PTX5000 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. 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-PTX-X8-64G is supported only on the new Control Board CB2-PTX.

- **New Control Board support (PTX5000)**—Starting with Release 15.1F4, Junos OS supports the Routing Engine RE-PTX-X8-64G with an enhanced Control Board (CB) on PTX5000 routers. The CB supports chassis management and 16 additional 10-Gigabit Ethernet ports with small form-factor pluggable plus transceivers (SFP+) on the front panel of the router to support multichassis applications.

The enhanced CB consists of the following components:

- Ethernet switch used for intermodule communication
 - PCI Express bus to connect to the Routing Engine
 - PCI Express switch to connect to the SIBs
 - Switch Processor Mezzanine Board (SPMB)
- **High capacity single-phase AC PDU (PTX5000)**—In Junos OS Release 15.1F3, a single-phase AC power distribution unit (PDU)—PDU2-PTX-AC-SP—is introduced to provide power to the PTX5000 chassis. The PDU provides a single-phase AC input connection from the customer's AC source, an I/O interface to the power supply modules (PSMs), and a DC power connection to the system midplane. The PDU is powered by either eight 30-A or eight 20-A single-phase sources. Each of the eight PSMs connected to the AC PDU receives single-phase input.

- **New horizontal fan tray FAN3-PTX-H (PTX5000)**—Starting in Junos OS Release 15.1F3, the FAN3-PTX-H horizontal fan tray is supported on PTX5000 routers.
- **New FPCs FPC3-PTX-U1-L, FPC3-PTX-U1-R, FPC3-PTX-U2-L, FPC3-PTX-U2-R, FPC3-PTX-U3-L, and FPC3-PTX-U3-R (PTX5000)**—Starting in Junos OS Release 15.1F3, the FPC3-PTX-U1-L, FPC3-PTX-U1-R, FPC3-PTX-U2-L, FPC3-PTX-U2-R, FPC3-PTX-U3-L, and FPC3-PTX-U3-R FPCs are supported on PTX5000 routers. The FPCs provide the following throughput:
 - FPC3-PTX-U1-L and FPC3-PTX-U1-R—1.0 Tbps
 - FPC3-PTX-U2-L and FPC3-PTX-U2-R—2.0 Tbps
 - FPC3-PTX-U3-L and FPC3-PTX-U3-R—3.0 Tbps

When installing these third-generation FPCs on the PTX5000 chassis, you must also install the following hardware:

- New SIB SIB3-PTX5K
- New horizontal fan tray FAN3-PTX-H

Some new features provided by these third-generation FPCs can be accessed only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

Some of the new features include the following:

- Filter-based generic routing encapsulation (GRE) for IPv4 and IPv6 tunneling uses firewall filters to provide decapsulation of GRE traffic. The filter-based GRE de-encapsulation also supports **routing-instance** as an action.
- **promote gre-key** statement for configuring gre-key as one of the matches in a filter.
- **gtp-tunnel-endpoint-identifier** statement for including hash calculation for IPv4 or IPv6 packets in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations.
- Longer configuration ranges for Bidirectional Forwarding Detection (BFD) protocol intervals.
- Enhanced support for up to two million routes per chassis.
- **New SIB SIB3-PTX5K (PTX5000)**—Starting in Junos OS Release 15.1F3, the SIB3-PTX5K SIB is supported on PTX5000 routers.
- **New PIC P3-24-U-QSFP28 (PTX5000)**—Starting in Junos OS Release 15.1F3, the PIC P3-24-U-QSFP28 is supported on PTX5000 routers. The P3-24-U-QSFP28 PIC has 24 ports configurable as either 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports.



NOTE: This PIC does not support 100-Gigabit Ethernet ports.

To install the P3-24-U-QSFP28 PIC, you must have a third-generation FPC installed on your system.

- **New P3-15-U-QSFP28 PIC (PTX5000)**—Starting in Junos OS Release 15.1F5, the PIC P3-15-U-QSFP28 is supported on PTX5000 routers that have third-generation FPCs installed.



NOTE: To install the P3-15-U-QSFP28 PIC, you must have a third-generation FPC installed on your system.

Following is the available port configuration for each FPC:

- FPC3-PTX-U1-L and FPC3-PTX-U1-R—10 ports configurable as 10-Gigabit Ethernet ports (using a 4x breakout cable), 40-Gigabit Ethernet ports, or 100-Gigabit Ethernet ports.
- FPC3-PTX-U2-L and FPC3-PTX-U2-R—10 ports configurable as 10-Gigabit Ethernet ports (using a 4x breakout cable), 40-Gigabit Ethernet ports, or 100-Gigabit Ethernet ports.
- FPC3-PTX-U3-L and FPC3-PTX-U3-R—15 ports configurable as 10-Gigabit Ethernet ports (using a 4x breakout cable), 40-Gigabit Ethernet ports, or 100-Gigabit Ethernet ports.

- **The P1-PTX-24-10G-W-SFPP PIC is supported on third-generation FPCs (PTX5000)**—Starting in Junos OS Release 15.1F5, the P1-PTX-24-10G-W-SFPP PIC is supported on PTX Series routers that have third-generation FPCs installed.
- **The P2-10G-40G-QSFPP PIC is supported on third-generation FPCs (PTX5000)**—Starting in Junos OS Release 15.1F5, the P2-10G-40G-QSFPP PIC is supported on PTX Series routers that have third-generation FPCs installed.
- **The P2-100GE-OTN PIC is supported on third-generation FPCs (PTX5000)**—Starting in Junos OS Release 15.1F5, the P2-100GE-OTN PIC is supported on PTX Series routers that have third-generation FPCs installed.
- **Upgrade of FPCs in an operational PTX5000**—Starting in Junos OS Release 15.1F5, you can upgrade the first-generation FPCs or second-generation FPCs to third-generation FPCs in an operational PTX5000.

You might need to upgrade the following components before you can upgrade the FPCs in a PTX5000:

- SIBs
- Fan tray
- Power distribution unit
- Power supply module

(See the *PTX5000 Packet Transport Router Hardware Guide*.)

- **The ability for third-generation FPCs to interoperate with first-generation and second-generation FPCs (PTX5000)**—Starting in Junos OS Release 15.1F5, when third-generation FPCs are installed on a chassis with first-generation and second-generation FPCs, the FPCs can interoperate with each other.



NOTE: For the third-generation FPCs to interoperate with the previous FPCs, the **enhanced-mode** statement cannot be configured on the chassis. Also, the third-generation FPCs can only provide the same functionality as the first-generation and second-generation FPCs. Any advanced features that third-generation FPCs might provide are disabled.

- **The P1-PTX-2-100G-WDM PIC is supported on third-generation FPCs (PTX5000)**—Starting in Junos OS Release 15.1F6, the P1-PTX-2-100G-WDM PIC is supported on PTX Series routers that have third-generation FPCs installed.
- **Integrated photonic line card (IPLC) (PTX3000)**—Starting in Junos OS Release 15.1F6, the PTX3000 can provide a fully integrated photonic line system for converged core and metro core packet optical networks running point-to-point and ring topologies. The following optical components are available for the PTX3000:
 - Integrated photonic line card (IPLC) base module—Provides the combined functionality of a 32-port reconfigurable optical add/drop multiplexer (ROADM), optical amplifier, optical equalizer, and optical channel monitor on a single card.
 - IPLC expansion module—Increases the channel capacity of the IPLC node to 64 channels.

The external optical inline amplifier (optical ILA) provides periodic amplification of the optical line signal to enable long-distance transmission.

To complete the optical solution, Juniper's integrated 100G Coherent transponders can be leveraged here, along with the IPLC, IPLC expansion module, optical ILA, and CSD to provide an end-to-end, fully managed packet optical solution.

You can configure, manage, and monitor the IPLC through Junos Space Connectivity Services Director 2.0.

Optionally, you can use the Junos OS CLI, or your SNMP management.

- **Third-generation FPCs (PTX3000)**—Starting in Junos OS Release 15.1F6, third-generation FPCs are supported on PTX3000 routers. FPC3-SFF-PTX-U1 FPCs (model numbers FPC3-SFF-PTX-U1-L and FPC3-SFF-PTX-U1-R) support 1.0 Tbps of throughput. FPC3-SFF-PTX-U0 FPCs (model numbers FPC3-SFF-PTX-U0-L and FPC3-SFF-PTX-U0-R) support 500 Gbps of throughput.

Third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) are supported only on a PTX3000 with SIB3-SFF-PTX SIBs. Third-generation FPCs and FPC-SFF-PTX-P1-A first-generation FPCs can interoperate with each other in the same system.

Some features provided by these third-generation FPCs can be accessed only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level. These features include the following:

- Filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling.
- **promote gre-key** statement for configuring gre-key as one of the matches in a filter.

- **gtp-tunnel-endpoint-identifier** statement for including hash calculation for IPv4 or IPv6 packets in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations.
- Wider configuration range for Bidirectional Forwarding Detection (BFD) protocol intervals.
- Support for Layer 3 VPNs. The **vrf-table-label** statement is supported.
- Support for destination class usage (DCU) and source class usage (SCU) accounting.
- Support for up to two million routes in the forwarding table.



NOTE: For third-generation FPCs to interoperate with FPC-SFF-PTX-P1-A first-generation FPCs, the **enhanced-mode** statement cannot be configured on the chassis. Any advanced features that third-generation FPCs may provide are disabled.

- **SIB3-SFF-PTX SIBs (PTX3000)**—Starting in Junos OS Release 15.1F6, SIB3-SFF-PTX SIBs are supported on PTX3000 routers. The SIB3-SFF-PTX SIBs are required with third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1). The SIB3-SFF-PTX SIBs also support FPC-SFF-PTX-P1-A first-generation FPCs—third-generation FPCs and FPC-SFF-PTX-P1-A first-generation FPCs can interoperate with each other in the same system.
- **P3-24-U-QSFP28 PIC supported on third-generation FPC (PTX3000)**—Starting in Junos OS Release 15.1F6, the P3-24-U-QSFP28 PIC is supported on FPC3-SFF-PTX-U1 FPCs on the PTX3000. The P3-24-U-QSFP28 PIC has 24 ports configurable as either 10-Gigabit Ethernet ports or 40-Gigabit Ethernet ports. To configure the port speed, use the following command:


```
[edit chassis]
user@host# set fpc slot-number pic pic-number port port-number port-speed (10G | 40G)
```
- **P1-PTX-24-10G-W-SFPP, P2-10G-40G-QSFPP, and P2-100GE-OTN PICs supported on third-generation FPCs (PTX3000)**—Starting in Junos OS Release 15.1F6, the P1-PTX-24-10G-W-SFPP, P2-10G-40G-QSFPP, and P2-100GE-OTN PICs are supported on third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) on the PTX3000.
- **Upgrading to third-generation FPCs and SIBs in an operational router (PTX3000)**—Starting in Junos OS Release 15.1F6, you can upgrade to third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) and SIB3-SFF-PTX SIBs in an operational PTX3000.
- **5-port 100-Gigabit DWDM OTN PIC with CFP2 (PTX3000 and PTX5000)**—Starting with Release 15.1F6, Junos OS supports the 5-port 100-Gigabit dense wavelength division multiplexing (DWDM) optical transport network (OTN) PIC (PTX-5-100G-WDM) with CFP2 analog coherent optical (CFP2-ACO) pluggable optics on third-generation FPCs on the PTX3000 and PTX5000 series routers. The 5-port 100-Gigabit DWDM OTN PIC supports the following features:

- Transparent transport of five 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.

General Routing

- **Support for virtualization on RE-PTX-X8-64G (PTX5000)**—Starting with Junos OS Release 15.1F3, the Routing Engine RE-PTX-X8-64G for PTX5000 supports virtualization.

Virtualization enables the router to support multiple instances of Junos OS and other operating systems on the same Routing Engine. However, for Junos OS Release 15.1F3, 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

Interfaces and Chassis

- **Support for including LOCAL-FAULT and REMOTE-FAULT information (PTX Series)**—Starting in Junos OS Release 15.1F3, PTX Series routers add the ability to display LOCAL-FAULT and REMOTE-FAULT information in the output of the **show interfaces et-fpc/pic/port** command.
- **Support for configuring chassis temperature thresholds (PTX Series)**—Starting in Junos OS Release 15.1F3, the **chassis [fpc| sib| cb]** statement is supported to define the thresholds at which the fans change speeds, the system is shut down, or an alarm is sent. The **chassis [fpc| sib| cb] threshold action to take** configuration statement is configured at the **[edit]** hierarchy level.
- **Support for configuring the port speed (PTX5000)**—Starting in Junos OS Release 15.1F3, the **port speed** configuration statement is used to configure the port speed on interface modules that support multiple port speeds. The **port-speed 10G | 40G | 100G** configuration statement is configured at the **[edit chassis fpc slot-number pic pic-number port port-number]** hierarchy level.

- **Support for configuring interface loopback (PTX5000)**—Starting in Junos OS Release 15.1F3, the **loopback (local | remote)** configuration statement is used to specify whether local or remote loopback is enabled. Specifying this information enables you to test the transceiver cable connection from the far end to the retimer interface without changing the cable. The **loopback (local | remote)** configuration statement is configured at the **[edit interfaces *interface-name* gigether-options]** hierarchy level.
- **Support for configuring the LED on a port to flash (PTX5000)**—Starting in Junos OS Release 15.1F3, the **led-beacon** command causes the LED for the specified port to flash green. This enables you to physically locate a specific optic port on the PIC. The **led-beacon** configuration statement is configured at the **[edit interfaces *interface-name* (with *port number*)]** hierarchy level.
- **Support for DCU and SCU accounting (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F5 for PTX5000 and Junos OS Release 15.1F6 for PTX3000, destination class usage (DCU) and source class usage (SCU) accounting are supported on routers that have third-generation FPCs installed.



NOTE: DCU and SCU accounting are supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Support for unicast RPF (PTX Series)**—Starting in Junos OS Release 15.1F6, you can configure unicast reverse path forwarding (RPF) to reduce the impact of denial of service (DoS) attacks on PTX Series routers that have third-generation FPCs installed.



NOTE: Unicast RPF is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Support for configuring port speed (PTX3000)**—Starting in Junos OS Release 15.1F6, the **port speed** configuration statement is used to configure the port speed on PICs that support multiple port speeds. The **port-speed (10G | 40G | 100G)** configuration statement is configured at the **[edit chassis fpc *slot-number* pic *pic-number* port *port-number*]** hierarchy level.
- **Support for configuring interface loopback (PTX3000)**—Starting in Junos OS Release 15.1F6, the **loopback (local | remote)** configuration statement is used to specify whether local or remote loopback is enabled. Specifying this information enables you to test the transceiver cable connection from the far end to the retimer interface without changing the cable. The **loopback (local | remote)** statement is configured at the **[edit interfaces *interface-name* gigether-options]** hierarchy level.
- **Support for configuring the LED on a port to flash (PTX3000)**—Starting in Junos OS Release 15.1F6, the **led-beacon** command causes the LED for the specified port to flash green. When the LED lights green, you can physically locate a specific optic port on the

PIC. You configure the **led-beacon** statement at the **[edit interfaces *interface-name* (with port number)]** hierarchy level.

- **Synchronous Ethernet clock synchronization on third-generation FPCs (PTX3000)**—Starting in Junos OS Release 15.1F6, Synchronous Ethernet clock synchronization is supported on third-generation FPCs (FPC3-SFF-PTX-U0 and FPC3-SFF-PTX-U1) on the PTX3000.

Management

- **Junos Telemetry Interface enhancements (PTX Series)**—Junos Telemetry Interface enables you to export telemetry data from supported interface hardware. Line card sensor data, such as interface events, are sent directly to configured collection points without involving polling. All parameters are configured at the **[edit services analytics]** hierarchy level. 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. 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 firewall filter statistics, include the **resource /junos/system/linecard/firewall/** statement at the **[edit services analytics sensor *sensor-name*]** hierarchy level. Only interfaces configured on FPC3 are supported. FPC2 is not supported.

MPLS

- **Egress peer engineering of service labels (BGP, MPLS) and egress peer protection for BGP-LU (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F4 for PTX5000 and Junos OS Release 15.1F6 for PTX3000, you can enable traffic engineering of service traffic, such as MPLS LSP traffic between autonomous systems (ASs), using BGP labeled unicast for optimum utilization of the advertised egress routes. You can specify one or more backup devices for the primary egress AS boundary router. Junos OS installs the backup path in addition to the primary path in the MPLS forwarding table, which enables MPLS fast reroute (FRR) when the primary link fails. It provides support for the FRR protection backup scheme to do an IP lookup to determine a new egress interface.
- **Support for IS-IS segment routing (PTX 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:
 - **no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.
 - **node-segment**—Enable source packet routing at all levels.

- **source-packet-routing**—Enable the source packet routing feature.
- **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.
- **Support for IPv6 tunneling over an MPLS-based IPv4 network (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F5 for PTX5000 and Junos OS Release 15.1F6 for PTX3000, IPv6 tunneling over an MPLS-based IPv4 network using IPv6 Provider Edge (6PE) is supported on routers that have third-generation FPCs installed.

Multicast

- **SAFI 129 NLRI compliance with RFC 6514 (PTX Series)**—Starting with Junos OS Release 15.1F2, the Network Layer Reachability Information (NLRI) format available for BGP VPN multicast is changing from the de facto format of Subsequent Address Family Identifier (SAFI) 128 to Subsequent Address Family Identifier (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]`.

Network Management and Monitoring

- **Support for accounting profiles (PTX Series)**—Starting in Junos OS Release 15.1F6, you can configure accounting profiles to collect data on PTX Series routers that have third-generation FPCs installed.



NOTE: Configuring accounting profiles is supported only when the `enhanced-mode` statement is configured at the `[edit chassis network-services]` hierarchy level.

Routing Policy and Firewall Filters

- **Support for firewall feature matching on gre-key (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F3 on PTX5000 and Junos OS Release 15.1F6 on PTX3000, the `promote gre-key` statement is supported to configure gre-key as one of the matches in a filter. When `promote gre-key` is configured and gre-key is used in any of the terms in a filter, the entire filter is compiled in a way that optimizes its performance for gre-key matching. The `promote gre-key` configuration statement is configured at the `[edit firewall family family-name filter filter-name]` hierarchy level.
- **Support for configuring the GTP-TEID field for GTP traffic (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F3 for PTX5000 and Junos OS Release 15.1F6 for PTX3000, the `gtp-tunnel-endpoint-identifier` statement is supported to configure the hash calculation of IPv4 or IPv6 packets that are included in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations. The `gtp-tunnel-endpoint-identifier` configuration statement is configured at the `[edit`

forwarding-options hash-key family inet layer-4] or [edit forwarding-options hash-key family inet6 layer-4] hierarchy level.

- **Support for the no-decrement-ttl tunneling attribute (PTX Series)**—Starting in Junos OS Release 15.1F6, you can configure the **no-decrement-ttl** tunneling attribute for filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling.



NOTE: The no-decrement-ttl tunneling attribute is supported only when the **enhanced-mode** statement is configured at the [edit chassis network-services] hierarchy level.

Routing Protocols

- **IS-IS purge originator identification TLV (PTX 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 the purge. This makes it easier to locate the origin of the purge and its cause.
- **Restricting LSP flooding over IS-IS interfaces (PTX 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 *f* 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.
- **BGP labeled unicast supports stack of labels (PTX 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 unicast label. 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.
- **Micro loop avoidance when IS-IS link fails (PTX Series)**—Beginning with Release 15.1F5, Junos OS allows a device to defer IS-IS route download when an IS-IS link fails—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 backup path is not impacted. In this case, traffic flow towards a converged path is deferred until all other nodes are converged.

- **System performance enhancements for rpd, Packet Forwarding Engine, and kernel (PTX 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-jflow (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F4 on PTX5000 and Junos OS Release 15.1F6 on PTX3000, you can use inline-jflow's export capabilities with IP Flow Information Export (IPFIX) to define a flow record template suitable for IPv4 or IPv6 traffic.
- **Enhancements were made to inline jflow to support large scale profiles (PTX Series with third-generation FPCs)**—Starting in Junos OS Release 15.1F6, the maximum flow sessions available increases to 800 thousand for 2 Tbps third-generation FPCs and 1.2 million for 3 Tbps third-generation FPCs.

Software-Defined Networking

- **Dynamic acquisition of network topology (PTX 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 (PTX 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 (PTX 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.
- **IGP-based topology discovery (PTX 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.
- **PCC delegation of auto-bandwidth and TE++ (PTX 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.

User Interface and Configuration

- **Monitoring, detecting, and taking action on degraded physical 100-Gigabit Ethernet links to minimize packet loss (PTX3000 and PTX5000)**—Starting with Junos OS Release 15.1F6, you can monitor physical link degradation (indicated by bit error rate (BER) threshold levels) on Ethernet interfaces, and take corrective actions if the BER threshold value drops to a value in the range of 10^{-13} to 10^{-5} .

The following new configurations have been introduced at the **[edit interfaces interface-name]** hierarchy level to support the physical link degrade monitoring and recovery feature on Junos OS:

- To monitor physical link degrade on Ethernet interfaces, configure the **link-degrade-monitor** statement.
- To configure the BER threshold value at which the corrective action must be triggered on or cleared from an interface, use the **link-degrade-monitor thresholds (set value | clear value)** statement.
- To configure the link degrade interval value, use the **link-degrade-monitor thresholds interval value** statement. The configured interval value determines the number of consecutive link degrade events that are considered before any corrective action is taken.
- To configure link degrade warning thresholds, use the **link-degrade-monitor thresholds (warning-set value | warning-clear value)** statement.
- To configure the link degrade action that is taken when the configured BER threshold level is reached, use the **link-degrade action media-based** statement.
- To configure the link degrade recovery options, use the **link-degrade recovery (auto interval value | manual)** statement.

You can view the link recovery status and the BER threshold values by using the **show interfaces interface-name** command.

VPNs

- **Support for Layer 3 VPN (PTX3000 and PTX5000)**—Starting in Junos OS Release 15.1F5 for PTX5000 and Junos OS Release 15.1F6 for PTX3000, Layer 3 VPN is supported on routers that have third-generation FPCs installed.



NOTE: Layer 3 VPN is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Flow-aware transport pseudowire for BGP L2VPN and BGP VPLS (PTX Series)**—Starting with Junos OS Release 15.1F2, the flow-aware transport (FAT) label that is supported for BGP-signaled pseudowires such as L2VPN and VPLS is configured only on the label edge routers (LERs). This causes the transit routers or label-switching routers (LSRs) to perform load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet

inspection of the payload. The FAT flow label can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for VPWS and VPLS pseudowires.

- Related Documentation**
- [Changes in Behavior and Syntax on page 149](#)
 - [Known Behavior on page 151](#)
 - [Known Issues on page 154](#)
 - [Resolved Issues on page 156](#)
 - [Documentation Updates](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 169](#)
 - [Product Compatibility on page 172](#)

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 in Junos OS Release 15.1F6 for the PTX Series.

- [Hardware on page 150](#)
- [Forwarding and Sampling on page 150](#)
- [General Routing on page 150](#)
- [IPv6 on page 150](#)
- [Network Management and Monitoring on page 150](#)
- [Routing Policy and Firewall Filters on page 151](#)

Hardware

- **Powering on offline FPCs (PTX5000)**—Beginning in Junos OS Release 15.1F3 offline FPCs do not come online during reboots or other power management events. To bring such an FPC online:
 1. Delete the **fpc fpc-slot power off** statement from the **[edit chassis]** hierarchy level, if that statement is configured, and commit the configuration.
 2. Either issue the **request chassis fpc online slot fpc-slot** operational-mode CLI command or press and hold the **FPC ONLINE/OFFLINE** button for about 5 seconds until the green **OK** LED next to the button lights steadily.

Forwarding and Sampling

- **Deprecation of disable option (PTX3000)**—Beginning in Junos OS 15.1F5, the **disable** option has been deprecated at the **forwarding-options sampling instance instance-name family (inet | inet6 | mpls)** hierarchy level on PTX3000 routers. When configured, the option does not take effect, so packets continue to be sampled. Instead of the **disable** option, use the **deactivate forwarding-options sampling instance instance-name family (inet | inet6 | mpls)** command to prevent sampling.

General Routing

- Starting in Junos OS Release 15.1F4, when a user uses the **request interface switchover** command to switch over the egress traffic from the secondary link, which is already down, to the primary link, the following error message is displayed:

error: ae4 operation not permitted since backup or primary down

IPv6

- **IPv6 addresses with padded zeros in MIC or MS-MPC system log messages (PTX Series)**—Starting with Junos OS Release 15.1F5, all system log messages originating from MIC or MS-MPC line cards displays padded zeros in IPv6 addresses to make it 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.

Network Management and Monitoring

- **New 64-bit counter of octets for interfaces (PTX 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 (PTX 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](#).]

Routing Policy and Firewall Filters

- **Support for logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol (PTX 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.

Related Documentation

- [New and Changed Features on page 135](#)
- [Known Behavior on page 151](#)
- [Known Issues on page 154](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates](#)
- [Migration, Upgrade, and Downgrade Instructions on page 169](#)
- [Product Compatibility on page 172](#)

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1F6 for the PTX Series.

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

- [General Routing on page 151](#)

General Routing

- The following **request** commands are not available for the Routing Engine RE-PTX-X8-64G on the PTX5000:

- **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 following **show** commands are modified for the Routing Engine RE-PTX-X8-64G:

- **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-PTX-X8-64G:

- **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-PTX-X8-64G:

- **edit vmhost**
- **edit system processes app-engine-virtual-machine-management-service**
- During deletion and restoration of scaled configurations on PTX5000, error messages related to next hops are displayed.
- The guest Junos OS and the host OS on the PTX5000 use different time zones. Therefore, there is a difference between the timestamps in the system log files of Junos OS and the host OS. As a workaround, you can calculate the current difference between the time zones used by Junos OS and the host OS and work with logs that show this difference in time.
- The configuration of the smartd process, which monitors the status of the disk on the host OS of the PTX5000, 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.
- The date and time zones are synchronized from the admin guest Junos OS to host OS on the PTX5000 and use same time zones. Therefore, there is no difference in the timestamp in system log files of Junos OS and 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 the master Routing Engine is running and stable and all the FPCs are in operational state. If the other Routing Engine gets powered off accidentally, issue the command **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.

Related Documentation

- [New and Changed Features on page 135](#)
- [Changes in Behavior and Syntax on page 149](#)
- [Known Issues on page 154](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates](#)

- [Migration, Upgrade, and Downgrade Instructions on page 169](#)
- [Product Compatibility on page 172](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1F6 for the PTX Series.

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

- [General Routing on page 154](#)
- [Integrated Photonic Line Card and Expansion Module on page 154](#)
- [MPLS on page 155](#)
- [Routing Protocols on page 155](#)

General Routing

- On FPC3, IPv4 and IPv6 port mirroring is not supported on the egress interface. [PR1081329](#)
- 100G mode for 96x10G/24x40G PIC is not supported. [PR1129888](#)
- On PTX3000 routers using FPC3-SFF-PTX-U1 FPC, removing the master Control Board from the chassis ungracefully can result in packet loss. To avoid packet loss, use the request chassis Routing Engine master switch command to change the master Routing Engine and Control Board to be the backup before removal. [PR1137194](#)
- Traffic may drop during Routing Engine switchover. [PR1164107](#)
- For PTX5000 platforms with Broadway-based FPCs, BFD V6 sessions with aggressive timers or when CPU load is high on the Routing Engine might flap, eg., during GRES or if less than 10 milisecond timer has been configured. This is applicable only for IPV6 BFD sessions on PTX Series platforms for Broadway-based FPCs. [PR1165719](#)

Integrated Photonic Line Card and Expansion Module

- The **show chassis fpc** command does not properly display the temperature of the IPLC expansion module. [PR1183222](#)
- The **show chassis environment** command does not display the temperature for the IPLC expansion module. [PR118232](#)
- Trap notifications for TCA alerts are not always sent properly on the IPLC module. [PR1185291](#)

MPLS

- This is a rare scenario, where LSP's bw is modified to maximum possible value of the link bw in one commit, upon commit some of the LSPs will be delayed to signal to the new bandwidth. For example, while the static bw of LSPs is modified by apply-group, let's say 500 LSPs each has 1m reservation is modified by apply-group configuration to 2m where total link bw is 1g. In this case some of the LSPs will be delayed to signal because of non-availability of bw. These LSPs will be signaled when they have sufficient bw to signal. This will only delay the process, and does not fail them to resignal with new bw. [PR1125323](#)

Routing Protocols

- BFD sessions (both PPMAN and TOE based sessions) using firewall daemon (dfwd) for steering Rx packets to Anchor Packet Forwarding Engine with NSR configured may flap post switchover. This is restricted to PTX Series platforms with FPC 3 equipped. The issue is seen only with BFD over AE and multi-hop BFD sessions. [PR1112218](#)
- When we have two paths for the same route, the route gets pointed to Unilist NH, which in turn gets pointed to two separate Unicast NHs. The route is determined by OSPF and we have BFD enabled on one of the paths which runs through an l2circuit path. When the link on the l2circuit gets cut, the link flap is informed by BFD as well as through OSPF LSAs. Ideally the BFD should inform the link down event before the OSPF LSA. But at the current situation, the OSPF LSAs update the current event a second before BFD. Due to this reason, we do get the route to be pointing to a new UNILIST NH with the weights swapped. But the Unicast NH for which the L3 link is down, gets added to the Unilist NH, the BFD assumes the link to be up, and hence updates the weights inappropriately and hence we do see traffic loss. Once the BFD link down event is processed at the OSPF protocol level, now the route points to only Unicast NH and we do see traffic flowing through the currently active link. The traffic outage would be hardly for less than a second during FRR. Also, this can be avoided if the BFD keepalive intervals are maintained around 50 ms with multiplier of 3 as opposed to 100 ms with a multiplier of 3. [PR1119253](#)
- On dual-Routing Engine platforms, in scaling scenario (e.g., there are 6 million routes on "old" master Routing Engine), if graceful Routing Engine switchover (GRES) or graceful restart (GR) is not enabled, the routing protocol process (rpd) may crash on the "old" master Routing Engine after performing Routing Engine switchover. As a workaround, if possible, rebooting the "old" master Routing Engine (new backup Routing Engine) after switchover could avoid the issue. [PR1128023](#)

Related Documentation

- [New and Changed Features on page 135](#)
- [Changes in Behavior and Syntax on page 149](#)
- [Known Behavior on page 151](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates](#)
- [Migration, Upgrade, and Downgrade Instructions on page 169](#)

- [Product Compatibility on page 172](#)

Resolved Issues

This section lists the issues fixed in Junos OS Release 15.1F6 for the PTX Series. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

- [Resolved Issues: 15.1F6 on page 156](#)
- [Resolved Issues: 15.1F5 on page 161](#)
- [Resolved Issues: 15.1F4 on page 163](#)
- [Resolved Issues: 15.1F3 on page 165](#)
- [Resolved Issues: 15.1F2 on page 167](#)

Resolved Issues: 15.1F6

- [Class of Service \(CoS\) on page 156](#)
- [General Routing on page 156](#)
- [Interfaces and Chassis on page 159](#)
- [MPLS on page 160](#)
- [Network Management and Monitoring on page 160](#)
- [Platform and Infrastructure on page 160](#)
- [Routing Protocols on page 161](#)

Class of Service (CoS)

- In case of member links of an AE (aggregated Ethernet) interface scatter over multiple Packet Forwarding Engines, if the FPC where member links of the AE interface reside gets reset or interface disable, there may be a dip in the output of SNMP walk on AE related queue MIB (such as jnxCosQstatTxedPkts). The behavior is intermittent and not seen every time. [PR1122343](#)
- 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](#)

General Routing

- Enhance error handling when TL-chip encounters KHT memory parity error. With the enhancement, when KHT memory parity error is detected, the content of the memory will be corrected using software shadow copy. [PR1001052](#)
- Sending 500G Packet Forwarding Engine line rate with small size packets (< 168B) without pre-classifier enabled in ingress buffering block (IPW) is not supported, which leads to loss of control packets resulting in routing instances flapping. [PR1093170](#)
- On PTX Series platforms, SIB temperature-threshold configuration cannot be changed for off-lined SIBs. [PR1097355](#)

- The "clear services accounting flow" command should not be used in 15.1F4 or 15.1F5 releases on inline J-Flow on PTX5000 router for PTX Series. This command is specific to J-Flow and is not supported in these releases. [PR1117181](#)
- In certain rare conditions, FPC VoQ will wedge which will drop packets on ingress Packet Forwarding Engine for PTX Series routers. Since the wedge cannot be reproduced, detection of wedge condition is introduced that alarm would be raised once the wedge condition is detected within 10 seconds. [PR1127958](#)
- When 2M IPv6 routes point to a single NH then IFD down or PIC offline of the PIC hosting the port related to the NH would crash the FPC. [PR1129183](#)
- When hold-time Up is configured on IFDs of PTX, SyncE line clocks can fail by restarting the associated FPCs. In a particular test scenario with relatively large hold-time Up 60,000 ms, SyncE line clock can fail by restarting the associated FPC, and it is never recovered until the IFD is dis/enabled. Both 10GbE and 100GbE interfaces indicate the same clock fail issue. [PR1130084](#)
- Link Fault signaling message are not logged on PTX Series Gladiator FPCs. [PR1132114](#)
- On PTX Series platforms with FPC3 equipped, the ifinfo (i.e., ifinfo is a short-live process that is created when the show interfaces command is issued via the Junos CLI) might crash during making changes on AE member link while running show interface related command on AE interface (such as make speed change (mixed<=.10G) with continuous retrieval of the LACP status using show lacp interfaces ae <n>). It will recover itself and there is no service impact. [PR1139409](#)
- 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](#)
- When smart configuration is configured and the user does a VMhost regular upgrade, user observes an error print NG_RE_VMM: "/usr/bin/vmsmartd.sh: Command not found." This can be ignored. This has no impact on upgrade. [PR1144881](#)
- Due to buffer size issue for FPC-SFF-PTX-P1-A (PTX3000) and FPC2-PTX-P1A (PTX5000) will be hitting "ISSU RECONNECT TIMEOUT" or "READY Message Without Reconnect" during unified ISSU (it will be hit if this fix not available in the from build). [PR1155936](#)
- With Junos OS Release 15.1F3 or 15.1F4, when a PTX Series router is performing penultimate-hop popping (PHP) and has an egress interface on FPC3, the MPLS EXP rewrite rule will not take effect on the exposed header after POP. [PR1160486](#)
- As per design, whenever configuration is committed in Hendricks, chassisd tries to bring the OFFLINE SIBs ONLINE. The SIB might not have power because of hardware issues (such as seen in this PR where the SIB couldn't be brought ONLINE because of a voltage rail failure). In such cases, software shouldn't try to ONLINE the SIB. This attempt is

resulting in the software state for this SIB getting stuck in "TRANSITION" state, thereby resulting in all future offline/online commands to not take effect on this SIB. Fixed software to check if there are any hard errors on that SIB before trying to bring it ONLINE. Fix comes in 14.2R7, 15.1F6, 15.1R4. [PR1161110](#)

- When per-VLAN shaping rate is changed on the router with FPC3 and FPC2, the "expr_cos_scheduler_map_unbind_ifl @ 1658: Failed to drain ifl" message is displayed and a short traffic outage may be observed. [PR1161306](#)
- TTL-1 (IP and MPLS) packets are not reported as Normal-Discards. [PR1162572](#)
- The description of the newer PTX5000 Control Board has changed from "PTX 5K CB" to "Control Board 2" to make the description consistent with existing Control Boards. [PR1163501](#)
- RPD core seen in krt_comp_rnhc_msg_set_af_info at `../../../../src/junos/usr.sbin/rpd/lib/krt/krt_comp_util.c:179` observed during feature integration test. [PR1163553](#)
- Control packets may get dropped when Packet Forwarding Engine is under heavy congestion. [PR1163759](#)
- A vmcore may be seen if there are 65000 pfstats requests in a short amount of time (70-90 seconds). [PR1163772](#)
- On a full box FPC3 chassis, FPC CPU utilization reaches 100% for a short period of time (between 30 seconds to 1 minute) during AE member links speed changes. [PR1163830](#)
- PTX FPC2 might core-dump after FPC reload or non-GRES Routing Engine mastership change due to memory corruption when processing IPC messages. We have enhanced the way DFWD for interface-specific filters is being handled. With this fix when the TLV decode has errors we would not continue processing and problem should no longer be seen. [PR1164055](#)
- FPC3 should not send the MTU exceeded MPLS packets to the Routing Engine. [PR1164145](#)
- In case that LSP Ping echo request rate is higher than the implemented rate-limit spec for MPLS-OAM application on PTX Series, some bursty LSP Ping packets get dropped and PTX Series cannot send LSP Ping echo reply for these dropped LSP Ping echo request. It causes that VCCV BFD session get destroyed by L2CKT-OAM client due to LSP Ping Reply timed out for L2CKT PW on the other remote endpoint that sent those LSP Ping echo requests. Since the remote endpoint stops sending BFD packets for those VCCV BFD sessions having become in AdminDown, the corresponding VCCV BFD session on PTX gets down as CtlExpire with Detect Timer Expiry event. The existed MPLS-OAM rate-limiter was increased to support a scaled MPLS-OAM scenario, 450 pps per Packet Forwarding Engine likely. [PR1164880](#)
- In mixed mode when the PVQ configuration is changed while there is traffic flowing, ifds may go down. [PR1167096](#)
- On PTX Series platforms, when a high priority clock source (bits-a) goes down, the clock status transits from "locked to bits-a" to "holdover" to "acquiring" to "locked to bits-b", and when the bits-a comes up, the clock status reverts from "locked to bits-b"

to "holdover" to "acquiring" to "locked to bits-a". In such scenarios, it should not drop by "holdover" and "acquiring" states as per the standard. [PR1168000](#)

- When stats are polled for tc-count, it was retaining stats read from multiple Packet Forwarding Engines when members of AE are spread across Packet Forwarding Engines of same FPC. Stats read from later Packet Forwarding Engines was overwriting the stats read from earlier Packet Forwarding Engines. [PR1170226](#)
- After booting up PTX Series, some setting is not set correctly with 15.1F5.15 and will trigger loopback wedge error message. [PR1171101](#)
- Traceroute will not work if a PTX5000 router-based line card is the ingress line card of a transit router (SWAP router) in 15.1F5.14. [PR1171119](#)
- In a very rare cases, multiple Routing Engine switchovers may result in SNGPMB crash. [PR1176094](#)
- When running interfaces on QSFP28 PIC in 40G or 100G mode, some of the interfaces the QSFP28 PIC may not come up after a system reboot. This issue does not impact interface running in 10G mode. [PR1176641](#)
- Gladiator: idle insert settings wrong on some SIBs, after do GRES for 3-4 times, then ungracefully power off/on SIBs. [PR1177652](#)
- "show t6e-pic 0 fec" on vty of fpc does not show xcvr type for QSFP+ optics. [PR1181430](#)

Interfaces and Chassis

- On PTX Series routers, TX optical threshold value is shown incorrectly for the interfaces in the PIC P1-PTX-2-100G-WDM. This PR will fix only the TX power issue reported in the 2x100G DWDM OTN PIC. [PR1084963](#)
- This command shows the Tx/Rx power/thresholds and alarms incorrectly. Here are the changes to address this PR: 1. The "-18 to -5 dBm" as defined in the spec refers to the per wavelength power, but the thresholds in the CLI are the cumulative power which can be received from multiple channels. So the change will be to keep the values as is and remove the "Laser" from the Receive power thresholds. 2. Remove the "laser output power xxxxx" thresholds as they do not have corresponding alarms. 3. Change the "laser output power alarms" to "Tx power xxx alarm/warn", these do not have corresponding thresholds. 4. Rename the "Laser rx power xxxx threshold" to "Rx power xxxx" 5. Rename "laser xxx power" to "xxx power". 6. Add the two los alarm/warn thresholds, which when configured will change the defaults for Rx low power warning and alarm. "set interfaces et-x/y/z optics-options los-alarm-threshold"/"set interfaces et-x/y/z optics-options los-warning-threshold". [PR1115135](#)
- On PTX Series platforms, whenever there is config change followed by a commit on 100GbE interfaces, the continuity-check (OAM) interval of 100 ms should be insufficient. The system keeps busy for longer than 100 ms in processing the configuration commit, so it can miss the continuity-check messages, then trigger Interface (OAM) flap in the same PIC. [PR1164329](#)
- If QSFP28-100GBASE-LR4/QSFP+-40G-LPBK PICs speed is configured at chassis hierarchy. DCD was not reading speed specified in (set chassis fpc <fpc> pic <pic> port <port> speed <speed>) and as a result, when IFDs created using this configuration

are added in AE bundle along with IFD of any other kind of pics, DCD used to give commit error. DCD was able to read speed for other IFDs in AE bundle and was not able to read speed of IFDs on QSFP28 PIC and hence used to complain about speed mismatch Commit error: Interface ae0 with child links of mixed speed but link-speed mixed is not configured. [PR1167780](#)

MPLS

- 14.1R6 - P2MP + NSR may observe ~100 ms loss during Routing Engine switchover in steady state with link protection enabled. This is not observed in Junos OS Releases 14.2 and later. The problem is undergoing further triage and is targeted to enter the 14.1R7 release window. [PR1095488](#)
- When automatic bandwidth is enabled on router acting as ingress and transit routers for scaled LSPs (for example, 4k ingress LSPs and 36k transit LSPs), "Max AvgBW util" update for ingress LSP will get delayed (for example, the statistics interval is set to 60s, but get updated almost 3+ minutes). It is because, even though the "no-transit-statistics" configuration statement is enabled, transit LSPs' statistics are still being considered for polling. This is noticed when there are large number of transit LSPs (such as 36k transit LSPs) present on the router. [PR1154212](#)
- The output of "show rsvp interfaces" was not being displayed in desired format. For higher number units, FPC in slot higher than 9 or channelized interfaces, output was being moved to the next line. That was harder to read and matching was not possible as part of the output was in another line, so nothing was returned in the CLI.
atajer@sangria-re1# run show rsvp interface | match "et-11/0/11|lo0" et-11/0/11:0.0 et-11/0/11:0.4000 et-11/0/11:1.0 et-11/0/11:2.0 et-11/0/11:3.0 lo0.0 Up 0 100% 0bps 0bps 0bps 0bps With the fix for this PR, no matter how many characters we have in the interface name, there will always be a space separating interface name and state. Formatting is not done using fixed length spaces anymore. State will not be aligned in a single column for all interfaces, but customer can see data using match conditions. Commit went into 15.1R4, 15.1F6 & 16.1x. atajer@sangria-re1# run show rsvp interface | match "et-11/0/11|lo0" et-11/0/11:0.0 Down 0 100% 10Gbps 10Gbps 0bps 0bps et-11/0/11:0.4000 Down 0 100% 10Gbps 10Gbps 0bps 0bps et-11/0/11:1.0 Down 0 100% 10Gbps 10Gbps 0bps 0bps et-11/0/11:2.0 Down 0 100% 10Gbps 10Gbps 0bps 0bps et-11/0/11:3.0 Down 0 100% 10Gbps 10Gbps 0bps 0bps lo0.0 Up 0 100% 0bps 0bps 0bps 0bps [PR1170790](#)

Network Management and Monitoring

- Eventd might run out of memory and crash because of excessive kernel logging. [PR1162722](#)

Platform and Infrastructure

- 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 does not handle an internal flag properly:
"KERNEL/PFE 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](#)
- With delta-export command enabled, "show|compare" output still shows after last successful commit. [PR1129577](#)

Routing Protocols

- 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](#)
- Even though no information is actually changed (all IS-IS adjacencies remain the same, etc.) when the IS-IS 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 two different fragments might cause issues for IS-IS "clients", CSPF for example. [PR1159482](#)

Resolved Issues: 15.1F5

- [General Routing on page 161](#)
- [Interfaces and Chassis on page 163](#)
- [MPLS on page 163](#)
- [Platform and Infrastructure on page 163](#)
- [Routing Protocols on page 163](#)

General Routing

- It is reported that on PTX Series platforms, when the firewall filter is configured on the loopback interface of the device, due to bad error handling or NULL pointer, all the FPCs on device may continuously crash and be unstable. Because the issue is not reproducible, the trigger of the issue is not clear. [PR996749](#)
- On PTX Series platform with FPC3, if the egress sampled packets that are less than one cell size (128 bytes including Ethernet header), the packets get chopped this host cannot receive the full IPv4/6 header. This issue might affect all the egress sampled packets sent to host (this can be syslog, log, or any other feature that samples the packet to host in the egress pipeline). [PR1100505](#)
- Interface status may not change accurately when FPC CPU usage is 100%. The problem will be seen if the following conditions are met on Gladiator. 1) Laser off/on are introduced from remote end in order to change the status of an interface on local end (Gladiator). 2) On local end (Gladiator), the FPC (with the interface for which status change is expected) CPU is 100% during this laser on/off event. An interface will not change its status from Down to Up or Up to Down triggered because of laser off/on from end, when the FPC CPU usage is 100%. The interface will correctly reflect its status as soon as FPC CPU usage drops below 100%. [PR1130920](#)

- On PTX platform with FPC3 equipped, the ifinfo (i.e., ifinfo is a short-live process that is created, when the show interfaces command is issued via the Junos CLI) might crash during making changes on AE member link while running show interface related command on AE interface (such as make speed change (mixed<=.10G) with continuous retrieval of the LACP status using show lacp interfaces ae <n>). It will recover itself and there is no service impact. [PR1139409](#)
- In 15.1F4 "show chassis environment" "Routing Engine 0 CPU" does not show Routing Engine CPU temperature as 15.1F4 does not have fix to make available CPU temperature in Junos OS from HOST. "Routing Engine 0 CPU" instead shows maximum of Routing Engine inlet and exhaust sensors reading. [PR1140187](#)
- After removing child link from AE bundle, the AE interface statistics in SNMP MIB might show a spike. [PR1140533](#)
- When an FPCs i2c register contents are reported in syslogs, the names of failed power rails are printed with the message. [PR1140556](#)
- IPv6 packet will be dropped on FPC3 of PTX Series platform if length of the IPv6 payload is smaller than 8 bytes. [PR1141384](#)
- On PTX Series platform with FPC3, the octets of IPv4 source and destination addresses in firewall log are listed reversely, it might affect troubleshooting. IPv6 log works fine. This is a minor issue, no other service impact. [PR1141495](#)
- On PTX Series platform with FPC3, the "disable" configuration statement under the hierarchy "forwarding-options sampling instance <instance name> family <inet>" does not take effect, so the packets are still getting sampled. As a workaround, please use "deactivate" command instead of "disable" command. [PR1142279](#)
- For In-Line J-Flow feature: When send traffic at line rate, sometime Inactive timeouts are incrementing, this issue will be fixed in 15.1F4-S1 and 15.1F5. [PR1142977](#)
- Packet drops are not reported as info cell drops in FPC3. [PR1145321](#)
- When exceeding the filter term limitation, the filter programming fails. When it fails it cleans up the failed filter. During the cleanup operation it tries to clean the FPC resources which were never actually allocated as the filter programming failed before the allocation. As a result it leads to assertion/FPC crash. [PR1148249](#)
- Counter defined with tc_cntr as substring will not appear in mibwalk. [PR1151884](#)
- Inline-Jflow on PTX FPC Type3 cards is now using the correct Domain Identifier ID in IPFIX exports based on FPC slot number. Previously, a value of 0x4 was always used. [PR1151955](#)
- In affected releases, when inline-jflow is used, certain debug information is logged by default. This is just a cosmetic issue and the fix will prevent these messages from being generated. [PR1152216](#)
- When per-VLAN shaping rate is changed on the router with FPC3 and FPC2, "expr_cos_scheduler_map_unbind_ifl @ 1658: Failed to drain ifl" messages is displayed and a short traffic outage may be observed. [PR1161306](#)
- TTL-1 (IP and MPLS) packets are not reported as Normal-Discards. [PR1162572](#)

- The description of the newer PTX5000 Control Board has changed from "PTX 5K CB" to "Control Board 2" to make the description consistent with existing Control Board. [PR1163501](#)
- Control packets may get dropped when Packet Forwarding Engine is under heavy congestion. [PR1163759](#)
- FPC3 should not send the MTU exceeded MPLS packets to the Routing Engine. [PR1164145](#)

Interfaces and Chassis

- On dual Routing Engine platforms, when adding the logical interfaces (IFLs) and committing, due to the device control process (dcd) on backup Routing Engine may 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 backup Routing Engine. [PR1014098](#)
- Release prior to 15.1F5 has buffer size issue for FPC-SFF-PTX-P1-A (PTX3000) and FPC2-PTX-P1A (PTX5000). Prior releases cannot use unified ISSU to upgrade to Junos OS Release 15.1F5. [PR1155936](#)

MPLS

- When Auto BW is enabled on Ingress Router, MaxAvgBW Util field will get delayed when large number of transit LSPs present though no-transit-statistics-polling. Though no-transit-statistics-polling configuration statement is enabled, still transit LSPs are being considered for polling. This is noticed when there are large number of transit LSPs present on the router. [PR1154212](#)

Platform and Infrastructure

- With delta-export command enabled, "show |compare" output still shows after last successful commit. [PR1129577](#)

Routing Protocols

- 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](#)

Resolved Issues: 15.1F4

- [Class of Service \(CoS\) on page 164](#)
- [General Routing on page 164](#)
- [Routing Protocols on page 165](#)

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](#)

General Routing

- Sending 500G Packet Forwarding Engine line rate with small size packets (< 168B) without pre-classifier enabled in ingress buffering block (IPW) is not supported, which leads to loss of control packets resulting in routing instances flapping. [PR1093170](#)
- On PTX5000 platform, show interfaces voq CLI command is not updating VOQ drops stats correctly. [PR1127725](#)
- In 15.1F3 RPD core can be seen on previous master after performing Routing Engine switchover. [PR1128023](#)
- On Next Generation PTX Series platform, FPC CPU spike might be observed if there is optic failure or mis-configuration, for example, one side is configured in 100G and another side in 10/40G. [PR1129057](#)
- On rare occasions, the PCI link on the SPMB CPU may not come up correctly due to a software error. This will prevent the SIBs from coming up properly if the SPMB happens to be Master. The problem can be resolved by restarting the SPMB using the command "request chassis spmb slot <id> restart". [PR1129203](#)
- On FPC3 equipped next-generation PTX5000 (model number: PTX5000BASE2) platform, packets drop would be seen when mac-address of the interface which has "ethernet-ccc" or "ethernet-tcc" encapsulation, is changed. As a workaround, deleting the "ethernet-ccc" (or "ethernet-tcc") encapsulation of the interface, making the mac-address change, doing the commit, then reverting back the "ethernet-ccc" encapsulation, and again doing the commit would avoid the issue. [PR1129641](#)
- When one of SIBs is offlined on PTX Series router, the performance might be affected and unable to reach line rate. [PR1129733](#)
- On PTX Series platform, when receiving large amount of TTL=1 MPLS packets on AE interface, the LACP packets might be affected and causing the AE interface to flap. [PR1129739](#)
- On Next Generation PTX Series platform with 100G interface, when we disable and enable the interface, in rare conditions, the link comes up and quickly goes down and then again comes up. This is not the expected behavior. [PR1132611](#)
- On PTX Series platforms while performing PIC offline/online or interface flap, FPC might generate a core file. [PR1132689](#)
- If the total combined PPS traffic exceeds a Packet Forwarding Engine limit, the Packet Forwarding Engine does not send out MAC-PAUSE Frame out of its interface indicating the downstream device to slow down. All control plane protocol packets are affected.

In this case, the LACP hello packet is not sent out, so the LACP might flap, the traffic forwarding will be affected. [PR1136038](#)

- PTX FPC3 - ifinfo core seen@#0 0x0808fa95 in pif_agg_traffic (ifl=0xffffd554) at ../../../../src/junos/usr.sbin/ifinfo/libifinfo/ifinfo_agg.c:935 during AE member link speed change (mixed<=.10G). [PR1139409](#)

Routing Protocols

- 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](#)

Resolved Issues: 15.1F3

- [General Routing on page 165](#)
- [Infrastructure on page 167](#)
- [Interfaces and Chassis on page 167](#)
- [MPLS on page 167](#)
- [Network Management and Monitoring on page 167](#)
- [Platform and Infrastructure on page 167](#)
- [Routing Protocols on page 167](#)
- [Software Installation and Upgrade on page 167](#)

General Routing

- When a switchover is done from one Routing Engine to the other, in graceful-switchover redundancy mode, there is a brief period early in the transition of the SIB to online state, during which unsolicited (not corresponding to an attempt by the CPU to access the SIB via PCIe) errors are received at the downstream PCIe port on the CB to the SIB. The fix is to mute the generation of such errors during this brief period of the switchover. [PR1068237](#)
- The configured buffer-size will not take effect until either "transmit-rate" or "excess-rate" is configured. [PR1072179](#)
- Tunable SFP+ optics are not supported on P1-PTX-24-10G-W-SFPP PIC in Junos OS Release 15.1R1. On Tunable Optics in this PIC, with Junos OS Release 15.1R1, the wavelength will not be configurable and the tunable parameters will not be correctly displayed in the CLI. [PR1081992](#)
- The FPC on PTX Series router might crash and reboot when the Packet Forwarding Engine is handling a fatal error. When the error happened, "TQCHIPO: Fatal error pqt_min_free_cnt is zero" log message is seen. [PR1084259](#)
- On PTX Series platform with external clock synchronization interface configured, when both BITS external clocks are disconnected at the same time, the 100GbE-LR4 FINISAR interface might flap. This link flap issue is narrowed down to the operation of data-path FIFO within CFP. When both the BITS clocks are disconnected, the reference clock jumps to "free-running" mode. This transition leads to a phase shift in reference clock.

Due to this phase shift, the data rates into and out of the FIFO will temporarily not match, leading to a FIFO over-run or under-run condition. This over-run or under-run condition forces a FIFO reset, and the output signal is distorted. So the far-end interface detects 'local-fault', then returns 'remote-fault' back to the near-end, hence a link flap. After change for this PR: - User needs to manually configure FPC recovered clock port for each clock put into "chassis synchronization source". Only one clock of each FPC can be put into "chassis synchronization source". [PR1091228](#)

- On PTX Series platform, if there are scaled configurations (for example, 5,000 routes and each of them with 64 ECMP paths configured) on a single interface and L2 rewrite profile is applied for the interface, the FPC may crash when deactivating and then activating the CoS configuration of the interface. [PR1096958](#)
- On PTX Series platform, SIB temperature-threshold configuration cannot be changed for off-lined SIBs. [PR1097355](#)
- When the PTX Series only has bits-a and bits-b as configured clock sources (and there is no interface on FPC configured as clock source), and it is losing signal from both of bits-a and bits-b simultaneously, clock sync state will go to FREERUN mode immediately. This is unexpected behavior. After the fix of this PR, clock sync state will stay HOLDOVER, then will go to FREERUN mode after the timeout. [PR1099516](#)
- Starting with Junos OS Release 14.1, Entropy Label Capability is enabled by-default on all Juniper Networks PTX Series routers. On PTX Series transit LSRs that carry LSPs with Entropy Label Capability, packet loss can be observed due to data errors when one or more labeled route entries are not properly removed from the hash table (i.e., following LSP optimization or MBB event) because the 'stale' entries are pointing to corrupted route memory. As a result, when the MPLS label that's associated with the 'stale' entry is re-used, data errors are seen for packets using the corresponding label. [PR1100637](#)
- On PTX Series platform, when yanking out FPC or SIB ungracefully (for example, pulling the line card out of the chassis unintentionally when the line card is carrying the traffic), there might be small probability that it can impact any of the FPCs with Grant Scheduler (GS) and Request Table (RT). Fatal interrupt occurred. [PR1105079](#)
- Suppose we have x MLDP LSPs and AE members with y child members. Then the scale supported is $4x*y < 60K$. [PR1107077](#)
- No decrement ttl does not work for incoming v6 traffic over MPLS IPv4 core. [PR1115203](#)

Infrastructure

- On PTX Series platform with Junos OS Release 15.1R1 and later, 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](#)

Interfaces and Chassis

- 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](#)

MPLS

- In the output of the CLI command "traceroute mpls ldp", the addresses of the interfaces on transit PTX Series routers might be shown as "127.0.0.1". [PR1081274](#)

Network Management and Monitoring

- Due to inappropriate cleanup in async library, disabling multiple interfaces while SNMP is polling interface oids might cause mib2d process to crash. [PR1097165](#).
- 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

- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)

Routing Protocols

- With any single-hop BFD session and MPLS OAM BFD session configured over the same interface, when the interface is disabled and enabled back immediately (e.g., a delay of 10 seconds between the two commit check-ins), the single-hop BFD session might get stuck into Init-Init state because the Down packet is received from other end for MPLS BFD session on the same interface might get demultiplexed to single-hop BFD session incorrectly. [PR1039149](#)

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](#)

Resolved Issues: 15.1F2

- [Forwarding and Sampling on page 168](#)
- [General Routing on page 168](#)

- [Interfaces and Chassis on page 168](#)
- [MPLS on page 168](#)

Forwarding and Sampling

- In PTX Series Carrier-Grade Service Engine (CSE) jflow solution environment, because the sampling process (sampled) may get into a continuous loop when handling asynchronous event (for example, aggregated tethered services interface flapping, or route update, or IFL/IFD update), the sampled process may never come out of that loop, which may result in high CPU usage (up to 90% sometimes). Also, the flabel might be exhausted because sampled is not able to consume states (such as route updates, interface updates) generated by kernel, and finally the router would not make any updates. [PR1092684](#)

General Routing

- Prior to this fix, "show interface diagnostics optics" command shows output for all four lanes for 10G ports of 48x10GE 12x40GE QSFP+ PIC. Normal behavior would be to display output for only the lane that the port belongs to. [PR959514](#)
- On PTX 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 unified in-service software upgrade (ISSU). The interrupt might be prevented after performing unified ISSU due to disabling the interrupt registers before unified ISSU, but never restored after. [PR1059098](#)
- The configured buffer-size will not take effect until either "transmit-rate" or "excess-rate" is configured. [PR1072179](#)
- On Junos OS Release 15.1R1, when the multicast next-hop is changed, the grafting and pruning operations take more time than before. [PR1090608](#)

Interfaces and Chassis

- If we load jinstall/jinstall64 image on PTX Series and if we have CFM configured over AE interfaces, this issue will be seen. [PR1085952](#)

MPLS

- 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](#)
- In Junos OS Release 14.1 and later, the "load-balance-label-capability" configuration statement is introduced to enable the router to push and pop the load-balancing label, which causes LDP and RSVP to advertise the entropy label TLV to neighboring routers. The PTX Series routers have the capability and it is reflected in their default forwarding-options configuration. However, there is a software defect in the way that Entropy Label Capability (ELC) TLV is encoded in the LDP label mapping message, which might cause the LDP session between the routers to go down. [PR1065338](#)

- Related Documentation**
- [New and Changed Features on page 135](#)
 - [Changes in Behavior and Syntax on page 149](#)
 - [Known Behavior on page 151](#)
 - [Known Issues on page 154](#)
 - [Documentation Updates](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 169](#)
 - [Product Compatibility on page 172](#)

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 PTX Series. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.



NOTE: When upgrading your PTX Series router to Junos OS 15.1F5, make sure that the interface device's (IFD) MTU limit is not set to less than 288 bytes. If the MTU is set to less than 288 bytes, your router will not upgrade properly.

- [Upgrading Using Unified ISSU on page 169](#)
- [Upgrading a Router with Redundant Routing Engines on page 169](#)
- [Basic Procedure for Upgrading to Release 15.1F5 on page 170](#)

Upgrading Using Unified ISSU



CAUTION: This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for PTX 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 in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [High Availability Feature Guide for Routing Devices](#).

Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Basic Procedure for Upgrading to Release 15.1F5

When upgrading or downgrading Junos OS, use the **jinstall** package. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#). Use other packages, such as the **jbundle** package, only when so instructed by a Juniper Networks support representative.



NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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



NOTE: 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 router, 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 for Routing Devices](#).



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.

The download and installation process for Junos OS Release 15.1F4 is different from previous Junos OS releases.

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



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.

The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release. Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes. Rebooting occurs only if the upgrade is successful.

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

```
user@host> request system software add validate reboot source/jinstall-15.1
F511-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-15.1  
F511-export-signed.tgz
```

Replace the ***source*** with one of the following values:

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

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

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 15.1F5 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.

Related Documentation

- [New and Changed Features on page 135](#)
- [Changes in Behavior and Syntax on page 149](#)
- [Known Behavior on page 151](#)
- [Known Issues on page 154](#)
- [Resolved Issues on page 156](#)
- [Documentation Updates](#)
- [Product Compatibility on page 172](#)

Product Compatibility

- [Hardware Compatibility on page 173](#)

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 PTX 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:

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

Related Documentation

- [New and Changed Features on page 135](#)
- [Changes in Behavior and Syntax on page 149](#)
- [Known Behavior on page 151](#)
- [Known Issues on page 154](#)
- [Resolved Issues on page 156](#)
- *Documentation Updates*
- [Migration, Upgrade, and Downgrade Instructions on page 169](#)

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:
<http://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:
<http://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:
<http://www.juniper.net/techpubs/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, 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 <http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <http://www.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <http://www.juniper.net/customers/support/>
- Search for known bugs: <http://www2.juniper.net/kb/>
- Find product documentation: <http://www.juniper.net/techpubs/>
- Find solutions and answer questions using our Knowledge Base: <http://kb.juniper.net/>
- Download the latest versions of software and review release notes: <http://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <http://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <http://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <http://www.juniper.net/cm/>

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

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.

- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, see <http://www.juniper.net/support/requesting-support.html>.

Revision History

23 March 2017—Revision 8, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

15 December 2016—Revision 7, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

27 October 2016—Revision 6, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

25 August 2016—Revision 5, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

18 August 2016—Revision 4, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

16 August 2016—Revision 3, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

14 July 2016—Revision 2, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

7 July 2016—Revision 1, Junos OS Release 15.1F6— MX Series, PTX Series and T Series.

Copyright © 2017, Juniper Networks, Inc. All rights reserved.

Juniper Networks, Junos, Steel-Belted Radius, NetScreen, and ScreenOS are registered trademarks of Juniper Networks, Inc. in the United States and other countries. The Juniper Networks Logo, the Junos logo, and JunosE are trademarks of Juniper Networks, Inc. All other trademarks, service marks, registered trademarks, or registered service marks are the property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.