

# Junos<sup>®</sup> OS Release 15.1F4 for the MX Series, PTX Series, and T Series

21 April 2016

## Contents

Introduction .....	4
Junos OS Release Notes for MX Series 3D Universal Edge Routers and T Series	
Core Routers .....	4
New and Changed Features .....	4
Hardware .....	5
General Routing .....	6
Interfaces and Chassis .....	6
Management .....	7
Multicast .....	7
Network Management and Monitoring .....	7
Platform and Infrastructure .....	7
Routing Protocols .....	8
Software Installation and Upgradation .....	9
Software-Defined Networking .....	9
Subscriber Management and Services (MX Series) .....	12
VPNs .....	12
Changes in Behavior and Syntax .....	12
MPLS .....	13
Network Management and Monitoring .....	13
Subscriber Management and Services (MX Series) .....	13
VPNs .....	14
Known Behavior .....	14
General Routing .....	14
Interfaces and Chassis .....	16
Known Issues .....	17
Forwarding and Sampling .....	17
General Routing .....	17
Interfaces and Chassis .....	19
Layer 2 Features .....	20
MPLS .....	20
Network Management and Monitoring .....	20
Platform and Infrastructure .....	21

Routing Protocols . . . . .	22
Services Applications . . . . .	22
User Interface and Configuration . . . . .	22
VPNs . . . . .	23
Resolved Issues . . . . .	23
Resolved Issues: 15.1F4 . . . . .	23
Resolved Issues: 15.1F3 . . . . .	29
Resolved Issues: 15.1F2 . . . . .	38
Documentation Updates . . . . .	48
Migration, Upgrade, and Downgrade Instructions . . . . .	48
Basic Procedure for Upgrading to Release 15.1F4 . . . . .	49
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x) . . . . .	51
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1) . . . . .	52
Upgrade and Downgrade Support Policy for Junos OS Releases . . . . .	54
Upgrading a Router with Redundant Routing Engines . . . . .	54
Upgrading the Software for a Routing Matrix . . . . .	55
Upgrading Using Unified ISSU . . . . .	56
Downgrading from Release 15.1 . . . . .	56
Product Compatibility . . . . .	57
Hardware Compatibility . . . . .	57
Junos OS Release Notes for PTX Series Packet Transport Routers . . . . .	58
New and Changed Features . . . . .	58
Hardware . . . . .	59
General Routing . . . . .	60
Interfaces and Chassis . . . . .	61
Management . . . . .	62
MPLS . . . . .	62
Routing Policy and Firewall Filters . . . . .	62
Routing Protocols . . . . .	63
Services Applications . . . . .	63
Software-Defined Networking . . . . .	63
Changes in Behavior and Syntax . . . . .	64
General Routing . . . . .	65
Network Management and Monitoring . . . . .	65
Known Behavior . . . . .	65
Known Issues . . . . .	65
General Routing . . . . .	66
Interfaces and Chassis . . . . .	67
Routing Protocols . . . . .	67
Resolved Issues . . . . .	67
Resolved Issues: 15.1F4 . . . . .	67
Resolved Issues: 15.1F3 . . . . .	69
Resolved Issues: 15.1F2 . . . . .	71
Documentation Updates . . . . .	73
Migration, Upgrade, and Downgrade Instructions . . . . .	73
Upgrading Using Unified ISSU . . . . .	73
Upgrading a Router with Redundant Routing Engines . . . . .	74
Basic Procedure for Upgrading to Release 15.1F4 . . . . .	74

---

Product Compatibility . . . . .	76
Hardware Compatibility . . . . .	77
Third-Party Components . . . . .	78
Finding More Information . . . . .	78
Documentation Feedback . . . . .	78
Requesting Technical Support . . . . .	79
Self-Help Online Tools and Resources . . . . .	79
Opening a Case with JTAC . . . . .	79
Revision History . . . . .	80

## 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 15.1F4 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 15.1F4 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.

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)
- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F4 for the MX Series and the T Series.

- [Hardware on page 5](#)
- [General Routing on page 6](#)
- [Interfaces and Chassis on page 6](#)
- [Management on page 7](#)
- [Multicast on page 7](#)

- [Network Management and Monitoring on page 7](#)
- [Platform and Infrastructure on page 7](#)
- [Routing Protocols on page 8](#)
- [Software Installation and Upgradation on page 9](#)
- [Software-Defined Networking on page 9](#)
- [Subscriber Management and Services \(MX Series\) on page 12](#)
- [VPNs on page 12](#)

## 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, you must download and install the Junos Continuity software package for Junos OS Release 15.1F4.

---

## 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 15.1F5 and later

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 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.

## Management

---

- **Router telemetry data for hardware and software (MX Series)**—Starting in Junos OS Release 15.1F3, you can configure MX Series routers to export telemetry data from supported interface hardware. Line card sensor data such as interface RSVP TE LSP events are sent directly to configured collection points without involving polling. You configure all parameters at the **[edit services analytics]** hierarchy level. You can configure the exact interfaces and LSPs for export statistics using regular expression resource filter matches. Supported MPC hardware on MX Series routers is MPC1 through MPC6E.

## Multicast

---

- **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.

## Platform and Infrastructure

---

- **Virtual MX Series router (vMX)**—Starting in Junos OS Release 15.1F3, you can deploy vMX routers on x86 servers. vMX supports most of the features available on MX Series

routers and allows you to leverage Junos OS to provide a quick and flexible deployment. vMX provides the following benefits:

- Optimizes carrier-grade routing for the x86 environment
- Simplifies operations by consistency with MX Series routers
- Introduces new services without reconfiguration of current infrastructure
- **Flow caching support on virtual MX Series router (vMX)**—Starting in Junos OS Release 15.1F4, you can enable flow caching on vMX routers for SR-IOV deployments. You enable flow caching by configuring the **performance-mode** option at the **[edit chassis fpc 0]** 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.

### Software Installation and Upgradation

- **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.
- **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.
- **Support of Internet draft draft-ietf-pce-stateful-pce-07 for the stateful PCC implementation (MX Series and T 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.

- **Support of Internet draft draft-crabbe-pce-pce-initiated-lsp-03 for the stateful PCE-initiated LSP implementation (MX Series and T 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.

- **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.

## Subscriber Management and Services (MX Series)

---

- **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.

## 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.

### Related Documentation

- [Changes in Behavior and Syntax on page 12](#)
- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## 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 15.1F4 for the MX Series and T Series.

- [MPLS on page 13](#)
- [Network Management and Monitoring on page 13](#)

- [Subscriber Management and Services \(MX Series\) on page 13](#)
- [VPNs on page 14](#)

## 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 (M Series, MX Series, and T 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.

## Subscriber Management and Services (MX Series)



**NOTE:** Although present in the code, the subscriber management features are not supported in Junos OS Release 15.1F4. Documentation for subscriber management features is included in the Junos OS Release 15.1 documentation set.

- **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.

## VPNs

---

- **Clear all Internet key exchange (IKE), traffic encryption key (TEK), key encryption key (KEK) 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 4](#)
- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1F4 for the MX Series and T Series.

- [General Routing on page 14](#)
- [Interfaces and Chassis on page 16](#)

### 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 PTX5000, MX240, MX480, MX960, MX2010, and MX2020, error messages related to next hops are displayed.
- The guest Junos OS and the host OS on the PTX5000, MX240, MX480, MX960, MX2010, and MX2020 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 PTX5000, 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.

---

## 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
```

### Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1F4 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 17](#)
- [General Routing on page 17](#)
- [Interfaces and Chassis on page 19](#)
- [Layer 2 Features on page 20](#)
- [MPLS on page 20](#)
- [Network Management and Monitoring on page 20](#)
- [Platform and Infrastructure on page 21](#)
- [Routing Protocols on page 22](#)
- [Services Applications on page 22](#)
- [User Interface and Configuration on page 22](#)
- [VPNs on page 23](#)

### Forwarding and Sampling

- When "shared-bandwidth-policer" is configured with aggregate Ethernet (AE), if there are filters configured on the logical interface family (IFF) of the AE interface, the FPC may crash upon rebooting (it is also seen when new FPC coming up) because the running thread is stuck at the association of the filter which is in the resolved state (this happens when the filter has not yet come down to the Packet Forwarding Engine whereas its association has already been reached). It is a timing issue in the above circumstance. However, it could be consistently reproduced when moving links from one aggregate Ethernet to another and then rebooting the FPC by scripts. As a workaround, if it is possible, the administrator could disable all the filter configurations and then bring up the line card. [PR1113915](#)

### General Routing

- If precision-timers and traceoptions are enabled for BGP, then both main-thread and precision-timers pthread try to rotate the same tracefile without taking any locks. As a result, all the status commands for rpd and krt may get timed out. [PR1044141](#)
- 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](#)
- From Junos OS Release 13.3, configuration changes like activating "fast-lookup-filter", adding or deleting "interface-specific" or any other filter property, adding or deleting any term, or changing any match condition in any term in the filter, which updates the firewall filter in a rare sequence might result in loss of DMEM and kernel memory resources or some free error messages. This issue impacts only MX Series-based line

card and occurs in rare cases. The following log messages might be observed from Junos OS Release 13.3 [LOG: Err] jnh\_free(10040): ERROR [FW/0]:1 Paddr 0x00400000, addr 0x0, part\_type 0 call\_stack 0x404906b4 0x418ecdbc 0x418ed2e0 0x418baca0 0x418becc8. And the following log messages might be observed from Junos 14.2 [LOG: Err] FREE ERR FW[0]: FW, 1 dw (1 blk) @ PA 0x7c63f00d addr 0x23f00d [PR1077338](#)

- The rpd process might crash on both master and backup Routing Engines when a routing instance is deleted from configuration if the routing instance is cleaned up before the interface delete is received from device control daemon (dcd). This is a rare timing issue. [PR1083655](#)
- There are entries for PEM in jnxFruEntry in VMX. It is not necessary and is cosmetic. [PR1094888](#)
- 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. [PR1095598](#)
- When route convergence occurred, the new gateway address is not updated correctly in inline-jflow route-record table (route-record table is used by sampling), and the sampling traffic forwarding might be affected, but normal routing would be not affected. [PR1097408](#)
- 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 kernel AE iffamily when subscribers login/logout. [1097824](#)
- Fragmenting a special host outbound IP packet with invalid IP header length (IP header length is greater than actual memory buffer packet header length), can trigger NULL mbuf accessing and dereferencing, which may lead to a kernel panic. [PR1102044](#)
- 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 CLI command below, re# set interfaces ms-3/0/0 services-options ignore-errors tcp. [PR1120079](#)
- MX Series router acting as L2TP access concentrator (LAC) may not recognize the MLPPP protocol field (0x003d) in the inbound PPP packet from customer premise equipment (CPE) and could disconnect the session not respecting idle-timeout. The traffic forwarding might be affected. [PR1123233](#)
- When GRES is enabled and EVPN is configured, the kernel crash may be seen during the Routing Engine mastership switchover. [PR1126195](#)
- When Junos OS devices use 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. 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 other vendor device which can map the configured 'port description' in the PortID TLV is used. In such case, Junos OS displays the neighbor's PortDescription TLV in the 'Port info' field, and if the peer sets 'port description' whose TLV length is longer than 33 byte(included), Junos OS is not able to accept the LLDP packets, then discards 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 are not deleted properly from the backup PE and causing the new DF election failure. The traffic forwarding will be affected. As a workaround, we should configure single primary loopback address and remove "router-id" configuration statement on both multi-homing PEs. [PR1126875](#)
- On MX Series routers with MS-MIC, changing configuration of sampling input parameters, such as "rate" under forwarding-options, is not reflected without restarting the MS-MIC. [PR1131227](#)
- CLI output of "clear services sessions" gives an impression to the user that 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](#)
- In certain Junos OS releases, making changes to the router-advertisement protocol stanza does not trigger new RAs to be sent after the commit. [PR1132345](#)
- MXVC-Same subnet VC-heartbeat polling failed to recover. [PR1136119](#)

## Interfaces and Chassis

- Chap Local-name default to 8 characters. Should be 32. [PR996760](#)
- Two redundant logical tunnels (rlt) interfaces are configured with configuration 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 following log can be seen on OTN capable pics after each commit, which indicates incorrect stats TLV setting. No service impact found. `/kernel: ge-1/1/0: Unknown TLV type 356 /kernel: ge-1/1/0: Unknown TLV type 361 /kernel: ge-1/1/0: get tlv ppfeid 0xe-0/2/0: get tlv ppfeid 0xe-0/3/0: get tlv ppfeid 0xe-1/2/0: get tlv ppfeid 0xe-1/3/0: get tlv ppfeid 0xe-2/0/0: get tlv ppfeid 0xe-2/1/0: get tlv ppfeid 0xe-2/2/0: get tlv ppfeid 0xe-2/3/0: get tlv ppfeid 0xe-5/1/0: get tlv ppfeid 0xe-5/1/1: get tlv ppfeid 0xe-5/1/2: get tlv ppfeid 0xe-5/1/3: get tlv ppfeid 0xe-5/1/4: get tlv ppfeid 0xe-5/1/5: get tlv ppfeid 0xe-5/1/6: get tlv ppfeid 0xe-5/1/7: get tlv ppfeid 0xe-5/1/8: get tlv ppfeid 0xe-5/1/9: get tlv ppfeid 0.` [PR1057594](#)
- On PB-20C12-ATM2-SMIR PIC, port 0 and port 1 are configured with clock source as external. If Loss of signal (LOS) is inserted on port 0, port 0 will be down. The expected behavior is clock being used from port 1. But in this case, port 0 down will result in port 1 flapping and reporting SONET phase lock loop (PLL) errors. [PR1098540](#)

- When an interface is added as part of an interface-set configuration, then there should be a valid configuration defined for that interface. [PR1128038](#)
- MX-VC specific behavior for SNMP walk of jnxOperating\* containers was divergent from physical MX Series router. Returned to vergence. [PR1136414](#)

---

### Layer 2 Features

- In a VPLS network that has a combination of automatic-site-id and fixed site-id together, if automatic-site-id is configured the first time and without deactivating/activating the VPLS protocol, it might not auto-recover when site id collision is detected, and the VPLS connections are stuck in LN/RN state. As a workaround, we can deactivate/activate the VPLS protocol. [PR1065952](#)
- There was bug in handling the code while redistributing the xmit and adj entries to pppman, when the interface entry was in pending distribution state. [PR1116741](#)
- In scenario that DHCP relay is used along with Virtual Extensible Local Area Network (VXLAN), if DHCP discover packet is received with the broadcast bit set via a VXLAN interface on MX Series platform (which is acting as DHCP relay), the OFFER back from the DHCP server will not be forwarded back to the client over the VXLAN interface. Unicast offers (that is, DHCP offer packet with unicast bit set) over VXLAN and both broadcast and unicast offers over native VLAN interfaces work fine. [PR1126909](#)
- When AE is core facing ifl in ldp-mesh vpls instance with local-switching in it, the traffic is looped back. [PR1138842](#)

---

### MPLS

- RSVP graceful restart does not function for LSPs that have a forwarding adjacency (FA) label-switched path (LSP) as a next hop. [PR60256](#)
- In scenario of egress-protection using stub-alias advertise mode where Point of Local Repair (PLR) use 'dynamic-rsvp-lsp' in LDP link protection, if protected PE get isolated, unexpected packet drops will be observed. [PR1030815](#)
- In MPLS scenarios, removing the "family mpls" configuration from an outgoing interface may cause inet and/or inet6 nexthops associated with that interface to unexpectedly transit to dead state. Even adding back "family mpls" cannot restore it. [PR1067915](#)
- These benign error messages can occur when bringing up new interfaces or line cards and can be ignored. The Junos OS code will be cleaned up to remove these messages in later code. [PR1136033](#)
- There is no entropy-label for LDP route in scenario of LDP tunneling across 1hop RSVP LSP + explicit-null RSVP label. Either remove LDP tunneling or RSVP explicit-null will resolve the issue as work-around. [PR1142357](#)

---

### Network Management and Monitoring

- When a firewall filter has one or more terms which have MX Series-only match condition or actions, such filters will not be listed during SNMP query. This behavior is seen typically after Routing Engine reboot/upgrade/master-ship switch. Restarting mib2d process will cause to learn these MX Series-only filters: cli > restart mib-process After

mib2d restart, SNMP mib walk of firewall OIDs will: - list all the OIDs corresponding this MX Series-only filter - count correctly as configured in the filter Now, despite the SNMP mib walk for firewall OIDs lists all OIDs and appropriate values, messages logs will report the following logs for every interface that has this TRIO-only filter applied.

```
router-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE:
get_counter_list: failed in reading counter names ae33.1009-i: 288 (No such file or
directory) router-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE:
get_counter_list: failed in reading counter names ae31.1004-i: 257 (No such file or
directory) router-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE:
get_counter_list: failed in reading counter names ae33.1010-i: 289 (No such file or
directory) router-re0 mib2d[4616]: %DAEMON-3-MIB2D_RTSLIB_READ_FAILURE:
get_counter_list: failed in reading counter names ae31.1004-i: 257 (No such file or
directory) The above 2 issues are addressed in this PR fix. PR988566
```

- Mib2d cores while trying to re-add a lag child into the internal DB since the entry is already present in the internal DB. Before adding the child link, mib2d does a lookup on the tree to know if the entry is not already there. However, this lookup returns no results, since the child link is part of snmp filter-interface configuration. [PR1039508](#)
- A merge conflict was incorrectly resolved by changing snmp trap value of jnxDomLaneNotifications to 26. The correct value will always be 25. [PR1145144](#)

## Platform and Infrastructure

- LMEM is an internal memory in LU/XL ASIC chip. It has private and shared regions for Packet Processing Engines. LMEM data errors are very rare events caused by environmental factors (this is not created by software). Due to a software defect, an error in the shared LMEM region will result in corruption of critical data structures of Packet Processing Engines that causes unpredictable communication of LU/XL ASIC chip with MQ/XM ASIC chip. These events will corrupt the state in MQ/XM and lead to a MQ/XM wedge. The MQ/XM wedge would cause fabric blackhole and finally reboot the line card. [PR1082932](#)
- With MX Series-based FPC, Load balance hash seed will be changed after unified ISSU. Since the hash seed value will be reverted to original value by rebooting FPC, there would be hash value inconsistency in the system which might introduce blackholing on multicast flavor traffic (mcast or BUM on vpls/l2-bridge). [PR1086286](#)
- On MX Series-based platform, when the type of the IPv6 traffic is non-TCP or non-UDP (for example, next header field is GRE or No Next Header for IPv6), if the traffic rate is high (for instance, higher than 3.5Mpps), the packet re-ordering may occur. [PR1098776](#)
- The kernel next-hop acknowledgement timeout maximum interval configured (krt-nexthop-ack-timeout) under the CLI hierarchy "routing-options forwarding-table" has been increased to 400 seconds to avoid performance issues with scaled subscribers. [PR1102346](#)
- Improved VTY commands to show internal JNH memory usage. [PR1103660](#)
- 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 "show configuration |display set" command to view the configuration. [PR1134117](#)

## Routing Protocols

---

- Continuous soft core-dump may be observed due to bgp-path-selection code. RPD forks a child and the child asserts to produce a core-dump. The problem is with route-ordering. And it is auto-corrected after collecting this soft-assert-coredump, without any impact to traffic/service. [PR815146](#)
- In rare cases, rpd may write a core file with signature "rt\_notbest\_sanity: Path selection failure on ..." The core is 'soft', which means there should be no impact to traffic or routing protocols. [PR946415](#)
- BFD session might reset on commit if version is configured. The adaptive RX interval gets set to 0 which results in the reset. A sample configuration of BFD version is as following: protocols { bgp { bfd-liveness-detection { version 1; minimum-interval 1000; transmit-interval { minimum-interval 1000; } } } } [PR1045037](#)

## Services Applications

---

- In some cases after unified ISSU upgrade/GRES switch/jl2tpd restarts, if the subscriber is terminated during the unified ISSU/GRES/restart process, jl2tpd may core. [PR1109447](#)

## User Interface and Configuration

---

- User needs to wait until the page is completely loaded before navigating away from the current page. [PR567756](#)
- Using the Internet Explorer 7 browser, while deleting a user from the Configure > System Properties > User Management > Users page on the J-Web interface, the system is not showing warning message, whereas in the Firefox browser error messages are shown. [PR595932](#)
- If you access the J-Web interface using the Microsoft Internet Web browser version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, on the Trace Options tab) even though the flags are not configured. As a workaround, use the Mozilla Firefox Web browser. [PR603669](#)
- On the J-Web interface, next hop column in Monitor > Routing > Route Information displays only the interface address and the corresponding IP address is missing. The title of the first column displays "static route address" instead of "Destination Address." [PR684552](#)
- On the J-Web interface, Configure > Routing > OSPF > Add > Interface Tab is showing only the following three interfaces by default: - pfh-0/0/0.16383 - lo0.0 - lo0.16385 To overcome this issue and to configure the desired interfaces to associated ospf area-range, perform the following operation on the CLI: - set protocols ospf area 10.1.2.5 area-range 12.25.0.0/16 - set protocols ospf area 10.1.2.5 interface fe-0/3/1 [PR814171](#)

- When entering "restart r" incomplete command in CLI, even though there are multiple options available, command "restart routing" is executed finally. It should throw an error like "error: invalid daemon: r". [PR1075746](#)
- When committing a configuration with very long as-path, in this case the as-path name is almost 12000 characters long, the committed process might crash. The committed process restart results in a minimal impact of system. [PR1119529](#)

### VPNs

- In L2VPN or VPLS scenario with Junos OS Release 14.2R4, after executing some negative operations, e.g., deactivate/active BGP and IGP, or restart FPCs, the rpd process might crash due to a NULL pointer access in code. [PR1104472](#)
- Source PE will send the duplicate multicast traffic over P-tunnel when Receiver PE advertises the different label for IPV4 and IPV6 in type 1 BGP route. [PR1128376](#)

### Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)
- [Known Behavior on page 14](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## Resolved Issues

This section lists the issues fixed in Junos OS Release 15.1F4 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: 15.1F4 on page 23](#)
- [Resolved Issues: 15.1F3 on page 29](#)
- [Resolved Issues: 15.1F2 on page 38](#)

### Resolved Issues: 15.1F4

- [Class of Service \(CoS\) on page 24](#)
- [Forwarding and Sampling on page 24](#)
- [General Routing on page 24](#)
- [Interfaces and Chassis on page 26](#)
- [Layer 2 Features on page 26](#)
- [MPLS on page 27](#)

- [Network Management and Monitoring on page 27](#)
- [Platform and Infrastructure on page 27](#)
- [Routing Protocols on page 28](#)
- [Services Applications on page 29](#)
- [VPNs on page 29](#)

### ***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 doesn't support EEC" should be moved from notice to debug level. [PR1020161](#)
- There is a remote loop back feature in 802.3ah standard, where one end can put remote end into remote-loopback mode by sending 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 platform when we configure the lfm remote-loopback we run into problem state. In problem state we see that the 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 peer router. [PR1046423](#)
- After executing CLI command "show route extensive", routing protocol process (rpd) may get into infinite loop and not respond anymore because the command may get executed a couple of times itself. In this situation, rpd high CPU utilization (running over 90% sometimes) might be seen on the device, and also the memory which used to store the command output would not be freed during those executions (in normal utilization, the memory uses about 160KB, but in problematic situation, it can swell to 3GB size), which would lead to rpd crash eventually after memory exhaustion. [PR1104090](#)
- When Bridge domain in PBB-EVPN Routing instance is modified to add/remove ISIDs BD can get stuck in destroyed state. This happens when ISIDs in the Bridge domain are changed from 1 to many or many to 1. This is only noticed during configuration changes or initial deployment. [PR1107625](#)

- In rare condition, after Routing Engine switchover, the MPC PIC might be offline, and some error messages might be seen. [PR1110590](#)
- On dual Routing Engine MX Series platform, the "xe" interfaces of any of the line cards below may flap during unified in-service software upgrade (ISSU) due to missing support. The flapping may not happen every time and the probability of occurrence would increase if more number of SFP+ (e.g. SFP+-10G-SR) are connected on the FPC. The affected line cards are, \* MIC3-3D-10XGE-SFPP \* MPC4E-3D-32XGE-SFPP, MPC4E-3D-2CGE-8XGE \* MPC5E-40G10G, MPC5EQ-40G10G \* MX2K-MIC6-24XE, MX2K-MIC6-24XE-OTN. [PR1118379](#)
- On MX240/MX480/MX960/MX2010/MX2020 products with MPC2E-3D-NG/MPC2E-3D-NG-Q/MPC3E-3D-NG/MPC3E-3D-NG-Q with MIC-3D-4XGE-XFP, IFD flap detection is much slower. It might lead to high FRR time, some traffic might be lost. [PR1122589](#)
- 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 is 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 Point-to-Point Protocol (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 29](#)
- [Infrastructure on page 33](#)
- [Interfaces and Chassis on page 34](#)
- [Layer 2 Features on page 34](#)
- [MPLS on page 35](#)
- [Network Management and Monitoring on page 35](#)
- [Platform and Infrastructure on page 35](#)
- [Routing Protocols on page 38](#)
- [Software Installation and Upgrade on page 38](#)
- [VPNs on page 38](#)

### General Routing

- In MX Virtual Chassis (MX-VC) environment, if the private local nexthops 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=URI>

```
<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 can not 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 (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+ can not 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 reset, 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 Trio 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=JSA10701=SIRT\\_1&actp=LIST](http://kb.juniper.net/InfoCenter/index?page=content=JSA10701=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 39](#)
- [General Routing on page 39](#)
- [Infrastructure on page 41](#)
- [Interfaces and Chassis on page 41](#)
- [Layer 2 Features on page 43](#)
- [MPLS on page 43](#)
- [Network Management and Monitoring on page 44](#)
- [Platform and Infrastructure on page 44](#)
- [Routing Protocols on page 46](#)

- [Services Applications on page 47](#)
- [Software Installation and Upgrade on page 47](#)
- [User Interface and Configuration on page 47](#)
- [VPNs on page 47](#)

### ***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](#)
- 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 Engines 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 enable 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/Router 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 on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)
- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 15.1F4 documentation for the MX Series and T Series.

**Related  
Documentation**

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)
- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)
- [Product Compatibility on page 57](#)

## Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the MX Series 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 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. The `request system software validate in-service-upgrade` command, which allows the detection of any compatibility issues before actually issuing the `request system software in-service-upgrade` command to initiate unified ISSU, is not supported in Junos OS Release 15.1 while upgrading from earlier Junos OS releases.

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

Platform	FreeBSD 6.1-based Junos OS	FreeBSD 10.x-based Junos OS
MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	YES	YES
T640, T1600, T4000, TX Matrix, TX Matrix Plus	YES	NO

- [Basic Procedure for Upgrading to Release 15.1F4 on page 49](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) on page 51](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 52](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 54](#)
- [Upgrading a Router with Redundant Routing Engines on page 54](#)
- [Upgrading the Software for a Routing Matrix on page 55](#)
- [Upgrading Using Unified ISSU on page 56](#)
- [Downgrading from Release 15.1 on page 56](#)

### [Basic Procedure for Upgrading to Release 15.1F4](#)

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 10.x)

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

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

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:  
<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 all the countries except Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia, use the following command:

For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-15.1F4.11.tgz
```

For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-15.1F4.11.tgz
```

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

For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-15.1F4.11-limited.tgz
```

For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-15.1F4.11-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.



**NOTE:** After you install a Junos OS Release 15.1F4 **jinstall** package, you cannot issue the **request system software rollback** command to return to the previously installed Junos OS (FreeBSD 6.1) software. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

---

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

---

Products impacted: 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-15.1F4.11-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-15.1F4.11-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](#).

---

### Downgrading from Release 15.1

---

To downgrade from Release 15.1 to another supported release, follow the procedure for upgrading, but replace the 15.1 **jinstall** package with one that corresponds to the appropriate release.

---



**NOTE:** You cannot downgrade more than three releases. For example, if your routing platform is running Junos OS Release 11.4, you can downgrade the software to Release 10.4 directly, but not to Release 10.3 or earlier; as a workaround, you can first downgrade to Release 10.4 and then downgrade to Release 10.3.

---

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

#### Related Documentation

- [New and Changed Features on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)

- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Product Compatibility on page 57](#)

## Product Compatibility

- [Hardware Compatibility on page 57](#)

### 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 on page 4](#)
- [Changes in Behavior and Syntax on page 12](#)
- [Known Behavior on page 14](#)
- [Known Issues on page 17](#)
- [Resolved Issues on page 23](#)
- [Documentation Updates on page 48](#)
- [Migration, Upgrade, and Downgrade Instructions on page 48](#)

## Junos OS Release Notes for PTX Series Packet Transport Routers

---

These release notes accompany Junos OS Release 15.1F4 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 58](#)
- [Changes in Behavior and Syntax on page 64](#)
- [Known Behavior on page 65](#)
- [Known Issues on page 65](#)
- [Resolved Issues on page 67](#)
- [Documentation Updates on page 73](#)
- [Migration, Upgrade, and Downgrade Instructions on page 73](#)
- [Product Compatibility on page 76](#)

### New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F4 for the PTX Series.

- [Hardware on page 59](#)
- [General Routing on page 60](#)
- [Interfaces and Chassis on page 61](#)
- [Management on page 62](#)
- [MPLS on page 62](#)
- [Routing Policy and Firewall Filters on page 62](#)
- [Routing Protocols on page 63](#)
- [Services Applications on page 63](#)
- [Software-Defined Networking on page 63](#)

## 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 decapsulation will also support 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.

---

## 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. This allows 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.

## Management

---

- **Router telemetry data for hardware and software (PTX Series)**—Starting in Junos OS Release 15.1F3, you can configure PTX Series routers to export telemetry data from supported interface hardware. Line card sensor data such as interface RSVP TE LSP events are sent directly to configured collection points without involving polling. All parameters are configured at the `[edit services analytics]` hierarchy level. You can configure the exact interfaces and LSPs for export statistics using regular expression resource filter matches. Supported FPC hardware on PTX Series routers is limited to FPC3. The PTX Series FPC2 hardware is not supported.

## MPLS

---

- **Egress peer engineering of service labels (BGP, MPLS) and egress peer protection for BGP-LU (PTX Series)**—Starting in Junos OS Release 15.1F4, 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 FRR protection backup scheme to do an IP lookup to determine a new egress interface.

## Routing Policy and Firewall Filters

---

- **Support for firewall feature matching on gre-key (PTX Series)**—Starting in Junos OS Release 15.1F3, 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 (PTX Series)**—Starting in Junos OS Release 15.1F3, 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.

## 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.

## Services Applications

- **Support for inline-jflow (PTX with third-generation FPCs)**—Starting in Junos OS Release 15.1F4, 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.

## 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.

**Related  
Documentation**

- [Changes in Behavior and Syntax on page 64](#)
- [Known Behavior on page 65](#)
- [Known Issues on page 65](#)
- [Known Issues on page 65](#)
- [Migration, Upgrade, and Downgrade Instructions on page 73](#)
- [Documentation Updates on page 73](#)
- [Product Compatibility on page 76](#)

## 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 15.1F4 for the PTX Series.

- [General Routing on page 65](#)
- [Network Management and Monitoring on page 65](#)

## 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**

## 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.

### Related Documentation

- [New and Changed Features on page 58](#)
- [Known Behavior on page 65](#)
- [Known Issues on page 65](#)
- [Resolved Issues on page 67](#)
- [Documentation Updates on page 73](#)
- [Migration, Upgrade, and Downgrade Instructions on page 73](#)
- [Product Compatibility on page 76](#)

## Known Behavior

There are no changes to the known behavior, system maximums, limitations in hardware and software in Junos OS Release 15.1F4 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.

### Related Documentation

- [New and Changed Features on page 58](#)
- [Changes in Behavior and Syntax on page 64](#)
- [Known Issues on page 65](#)
- [Known Issues on page 65](#)
- [Documentation Updates on page 73](#)
- [Migration, Upgrade, and Downgrade Instructions on page 73](#)
- [Product Compatibility on page 76](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1F4 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 66](#)
- [Interfaces and Chassis on page 67](#)
- [Routing Protocols on page 67](#)

### General Routing

---

- PTX Series does not support the queuing PICs but by default Junos OS will program chassis scheduler map which will generate the following logs: "fpc2 COS(cos\_chassis\_scheduler\_pre\_add\_action:2140): chassis scheduler ipc received for non qpvc ifd et-2/1/3 with index 131 /kernel: GENCFG: op 8 (COS BLOB) failed; err 5 (Invalid)Fix: Adding check to stop sending chassis scheduler map on PTX platform." Fix: Adding check to stop sending chassis scheduler map on PTX Series platform. [PR910985](#)
- 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 the device may continuously crash and become unstable. Because the issue is not reproducible, the trigger of the issue is not clear. [PR996749](#)
- 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 OSPF protocol level, now the route points to only UNICAST NH and hence 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 as 3 as opposed to 100ms with a multiplier of 3. [PR1119253](#)
- If a static route points to a set of interfaces effectively resulting in static route pointing to a unilist nexthop, it is possible that the selector weights may not be initialized correctly resulting in traffic drop. You can mitigate this issue by deactivating and then activating the static route configuration. [PR1120370](#)
- 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](#)

- Using the "write core dump" vty command on FPC causes crash after the core is uploaded. [PR1139370](#)
- For In-Line JFlow feature: When send traffic at line rate, sometimes Inactive timeouts are incrementing. [PR1142977](#)

### Interfaces and Chassis

- On PTX Series routers, TX optical threshold value is shown incorrect 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](#)
- On PTX platform "cfp\_lh\_update\_1sec\_pm\_var received" messages are periodically logged with Warning level. The severity of this message has been revised. [PR1089592](#)

### Routing Protocols

- On shmlog unsupported platforms (for example, PTX Series and T Series platforms), the following message might be seen after a configuration change: PTX-re0 rpd[42030] shmlog not initialized for PIM - not provisioned in platform manifest file The message does not indicate an error, it just indicates that shmlog is not supported on the PTX Series platform. The severity of the log has been reduced to INFO. [PR1065055](#)

#### Related Documentation

- [New and Changed Features on page 58](#)
- [Changes in Behavior and Syntax on page 64](#)
- [Known Behavior on page 65](#)
- [Known Issues on page 65](#)
- [Documentation Updates on page 73](#)
- [Migration, Upgrade, and Downgrade Instructions on page 73](#)
- [Product Compatibility on page 76](#)

## Resolved Issues

This section lists the issues fixed in Junos OS Release 15.1F4 for the PTX Series. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

- [Resolved Issues: 15.1F4 on page 67](#)
- [Resolved Issues: 15.1F3 on page 69](#)
- [Resolved Issues: 15.1F2 on page 71](#)

### Resolved Issues: 15.1F4

- [Class of Service \(CoS\) on page 68](#)
- [General Routing on page 68](#)
- [Routing Protocols on page 69](#)

### ***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, do the commit, then reverting back the "ethernet-ccc" encapsulation, and again do 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 condition, 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 exceed 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. And 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 69](#)
- [Infrastructure on page 71](#)
- [Interfaces and Chassis on page 71](#)
- [MPLS on page 71](#)
- [Network Management and Monitoring on page 71](#)
- [Platform and Infrastructure on page 71](#)
- [Routing Protocols on page 71](#)
- [Software Installation and Upgrade on page 71](#)

### ***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, "TQCHPO: 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, 5000 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 (ie. 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 below 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 72](#)
- [General Routing on page 72](#)

- [Interfaces and Chassis on page 72](#)
- [MPLS on page 72](#)

### ***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 58](#)
  - [Changes in Behavior and Syntax on page 64](#)
  - [Known Behavior on page 65](#)
  - [Known Issues on page 65](#)
  - [Documentation Updates on page 73](#)
  - [Migration, Upgrade, and Downgrade Instructions on page 73](#)
  - [Product Compatibility on page 76](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 15.1F4 documentation for the PTX Series.

- Related Documentation**
- [New and Changed Features on page 58](#)
  - [Changes in Behavior and Syntax on page 64](#)
  - [Known Behavior on page 65](#)
  - [Known Issues on page 65](#)
  - [Known Issues on page 65](#)
  - [Migration, Upgrade, and Downgrade Instructions on page 73](#)
  - [Product Compatibility on page 76](#)

## 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.

- [Upgrading Using Unified ISSU on page 73](#)
- [Upgrading a Router with Redundant Routing Engines on page 74](#)
- [Basic Procedure for Upgrading to Release 15.1F4 on page 74](#)

### 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 15.1F4

---

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
F411-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-15.1  
F411-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 15.1F4 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.

---

#### Related Documentation

- [New and Changed Features on page 58](#)
- [Changes in Behavior and Syntax on page 64](#)
- [Known Behavior on page 65](#)
- [Known Issues on page 65](#)
- [Known Issues on page 65](#)
- [Documentation Updates on page 73](#)
- [Product Compatibility on page 76](#)

## Product Compatibility

- [Hardware Compatibility on page 77](#)

## 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 58](#)
- [Changes in Behavior and Syntax on page 64](#)
- [Known Behavior on page 65](#)
- [Known Issues on page 65](#)
- [Known Issues on page 65](#)
- [Documentation Updates on page 73](#)
- [Migration, Upgrade, and Downgrade Instructions on page 73](#)

## 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](mailto:techpubs-comments@juniper.net), or fill out the documentation feedback form at <https://www.juniper.net/cgi-bin/docbugreport/>. If you are using e-mail, be sure to include the following information with your comments:

- Document or topic name
- URL or page number
- Software release version (if applicable)

## Requesting Technical Support

---

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at <http://www.juniper.net/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 located at <https://tools.juniper.net/SerialNumberEntitlementSearch/>.

## Opening a Case with JTAC

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

- Use the Case Management tool in the CSC at <http://www.juniper.net/cm/>.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to **support@juniper.net**. For documentation issues, fill out the bug report form located at <https://www.juniper.net/cgi-bin/docbugreport/>.

---

## Revision History

21 April 2016—Revision 8, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

24 March 2016—Revision 7, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

17 March 2016—Revision 6, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

3 March 2016—Revision 5, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

22 January 2016—Revision 4, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

21 January 2016—Revision 3, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

7 January 2016—Revision 2, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

28 December 2015—Revision 1, Junos OS Release 15.1F4— MX Series, PTX Series, and T Series.

Copyright © 2016, 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.