

# Junos<sup>®</sup> OS Release 15.1F2 for the M Series, MX Series, PTX Series, and T Series

7 January 2016

## Contents

Introduction	3
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers	3
New and Changed Features	3
High Availability (HA) and Resiliency	4
Interfaces and Chassis	4
MPLS	5
Multicast	5
Services Applications	7
Software Installation and Upgrade	7
System Management	7
User Interface and Configuration	7
VPNs	9
Changes in Behavior and Syntax	9
High Availability and Resiliency	9
Known Behavior	10
Known Issues	10
Class of Service (CoS)	11
Forwarding and Sampling	11
General Routing	11
Interfaces and Chassis	13
Layer 2 Features	14
MPLS	14
Network Management and Monitoring	15
Platform and Infrastructure	15
Routing Protocols	16
Services Applications	16
Subscriber Access Management	16
User Interface and Configuration	17
VPNs	17
Resolved Issues	18
Resolved Issues: 15.1F2	18

Documentation Updates . . . . .	28
Migration, Upgrade, and Downgrade Instructions . . . . .	28
Basic Procedure for Upgrading to Release 15.1F2 . . . . .	29
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x) . . . . .	31
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1) . . . . .	32
Upgrade and Downgrade Support Policy for Junos OS Releases . . . . .	34
Upgrading a Router with Redundant Routing Engines . . . . .	34
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 . . . . .	34
Upgrading the Software for a Routing Matrix . . . . .	36
Upgrading Using Unified ISSU . . . . .	37
Downgrading from Release 15.1 . . . . .	37
Product Compatibility . . . . .	38
Hardware Compatibility . . . . .	38
Junos OS Release Notes for PTX Series Packet Transport Routers . . . . .	39
New and Changed Features . . . . .	39
Class of Service (CoS) . . . . .	39
Multicast . . . . .	40
VPNs . . . . .	40
Changes in Behavior and Syntax . . . . .	40
Known Behavior . . . . .	41
Known Issues . . . . .	41
General Routing . . . . .	41
Interfaces and Chassis . . . . .	41
MPLS . . . . .	42
Routing Protocols . . . . .	42
Resolved Issues . . . . .	42
Resolved Issues: 15.1F2 . . . . .	42
Documentation Updates . . . . .	44
Migration, Upgrade, and Downgrade Instructions . . . . .	44
Upgrading Using Unified ISSU . . . . .	44
Upgrading a Router with Redundant Routing Engines . . . . .	45
Basic Procedure for Upgrading to Release 15.1F2 . . . . .	45
Product Compatibility . . . . .	47
Hardware Compatibility . . . . .	48
Third-Party Components . . . . .	49
Finding More Information . . . . .	49
Documentation Feedback . . . . .	49
Requesting Technical Support . . . . .	50
Self-Help Online Tools and Resources . . . . .	50
Opening a Case with JTAC . . . . .	50
Revision History . . . . .	51

---

## 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 15.1F2 for the M Series, 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 M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers

---

These release notes accompany Junos OS Release 15.1F2 for the M Series, 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 M Series, MX Series, and T Series routers. We recommend that you not use unified ISSU to upgrade from an earlier Junos OS release to Junos OS Release 15.1.

- 
- [New and Changed Features on page 3](#)
  - [Changes in Behavior and Syntax on page 9](#)
  - [Known Behavior on page 10](#)
  - [Known Issues on page 10](#)
  - [Resolved Issues on page 18](#)
  - [Documentation Updates on page 28](#)
  - [Migration, Upgrade, and Downgrade Instructions on page 28](#)
  - [Product Compatibility on page 38](#)

## New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F2 for the M Series, MX Series, and T Series.

- [High Availability \(HA\) and Resiliency on page 4](#)
- [Interfaces and Chassis on page 4](#)
- [MPLS on page 5](#)
- [Multicast on page 5](#)
- [Services Applications on page 7](#)

- [Software Installation and Upgrade on page 7](#)
- [System Management on page 7](#)
- [User Interface and Configuration on page 7](#)
- [VPNs on page 9](#)

---

## High Availability (HA) and Resiliency

- **Support for unified ISSU on MX Series routers with MPC5E and MPC6E (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Release 15.1F2, Junos OS supports unified in-service software upgrade (ISSU) on MX Series routers with MPC5E (MPC5E-40G10G, MPC5E-100G10G), MPC5EQ (MPC5EQ-40G10G, MPC5EQ-100G10G), and MPC6E (MX2K-MPC6E). Junos OS also extends support for unified ISSU on the following MICs that are supported on MPC6E:
  - *10-Gigabit Ethernet MIC with SFP+ (24 Ports)*
  - *10-Gigabit Ethernet DWDM OTN MIC (non-OTN mode only)*
  - *100-Gigabit Ethernet MIC with CFP2 (non-OTN mode only)*
  - *100-Gigabit Ethernet MIC with CXP (4 Ports)*

---

## Interfaces and Chassis

- **Support for ITU-T Y.1731 ETH-LM, ETH-SLM, and ETH-DM on aggregated Ethernet interfaces (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1F2, you can configure ITU-T Y.1731 standard-compliant Ethernet loss measurement (ETH-LM), Ethernet synthetic loss measurement (ETH-SLM), and Ethernet delay measurement (ETH-DM) capabilities on aggregated Ethernet (ae) interfaces. These performance monitoring functionalities are supported on MX Series routers with MPCs, where the same level of support for the Ethernet services OAM mechanisms as the level of support on non-aggregated Ethernet interfaces is available on aggregated Ethernet interfaces. ETH-DM is supported on MPC3E and MPC4E modules with only software timestamping. ETH-SLM is supported on MPC3E and MPC4E modules.
- **Routing Engine failover detection (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1F2, you use the **on-re-to-fpc-stale** configuration statement at the **[edit chassis redundancy failover]** hierarchy level to instruct the backup Routing Engine to take the mastership if the em0 interface fails on the master Routing Engine.
- **Support for targeted aggregated Ethernet distribution (MX Series routers with MPCs/MICs)**—In Junos OS Release 15.1F2 and later, you can direct traffic through specified links of a logical interface of an aggregate Ethernet bundle that is configured without link protection. This feature is supported on interfaces configured on MX Series MPCs and MICs.

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

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



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

## MPLS

- **Leaking MPLS routes to nondefault routing instances (MX Series routers with MPC/MIC interfaces)**—Starting in Junos OS Release 15.1F2, you can use the `import-labeled-routes` statement at the `[edit routing-instances routing-instance-name protocols vpls]` hierarchy level to specify one or more nondefault routing instances where you want MPLS pseudowire labeled routes to be leaked from the `mpls.0` path routing table in the master routing instance.

This capability prevents traffic loss in an L2VPN/VPLS configuration where the remote PE router is learned from the IGP in a nondefault routing instance. Because ingress-labeled routes are installed only in the master `mpls.0` table by default, no route is found in the `routing-instance-name.mpls.0` table when L2VPN/VPLS traffic received on the core-facing interface, and that traffic is dropped.

## Multicast

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

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

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

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

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

- **Improved scaling for multicast OIFs (MX Series)**—Starting with Junos OS Release 15.1F2, for both Rosen and next-generation MVPN, improvements have been made to increase the number of possible outgoing interfaces (OIFs) used in virtual routing and forwarding (VRF). Changes have also been made to improve the efficiency and scalability of PIM Join/Prune message processing.

These changes are implemented by default and do not need to be explicitly enabled. The following operational commands support the increased scaling.

- **show multicast next-hops terse**
  - **show multicast route oif-count**
  - **show multicast statistics interface**
  - **show pim join downstream-count**
- **Improved multicast convergence and RPT-SPT support for BGP-MVPN (MX Series)**—Starting with Junos OS Release 15.1F2, support for multicast forwarding-cache threshold is extended to rendezvous-point tree shortest-path tree (RPT-SPT) mode for BGP-MVPN. In addition, for both Rosen and next-generation MVPNs, PE routers across all sites should see the same set of multicast routes even if the configured forwarding-cache limit is exceeded.

To configure a specific threshold for MVPN RPT, set one or both of the **mvpn-rpt-suppress** and **mvpn-rpt-reuse** statements at the **[edit routing-instances *name* routing-options multicast forwarding-cache]** or **[edit logical system *name* routing-instances *name* routing-options multicast forwarding-cache]** hierarchy level.

In addition, the **show multicast forwarding-cache statistics** command provides information about both the general and RPT-suppression states. Likewise, a list of suppressed customer-multicast states can be seen by running the **show mvpn suppressed *general|mvpn-rpt inet|inet6* instance *name* summary** command.

## Services Applications

---

- **Support for inline MPLS Junos Traffic Vision with IPFIX and v9 (MX Series)**—Starting in Junos OS Release 15.1F2, support of the MX Series routers for the inline Junos Traffic Vision feature is extended to the MPLS family (MPLS and MPLS-IPv4 templates) consisting of the IP Flow Information Export (IPFIX) protocol and flow monitoring version 9 (v9). In previous releases, the inline Junos Traffic Vision feature is supported only for IPv4, IPv6, and VPLS families.

Inline Junos Traffic Vision feature is extended to MPC5E and MPC6E (XL-based chips) for VPLS address family. Also, Inline Junos Traffic Vision support using version 9 templates is extended to VPLS family.

## Software Installation and Upgrade

---

- **Validate system software against running configuration on remote host**—Beginning with Junos OS Release 15.1F2, you can use the `on (host host <username username> | routing-engine routing-engine)` option with the `request system software validate package-name` command to verify candidate system software against the running configuration on the specified remote host or Routing Engine.

## System Management

---

- **Statement introduced to deny hidden commands**—Starting in Release 15.1F2, Junos OS allows users to deny hidden commands to all users except root. To deny hidden commands to all users except root, use the `set system no-hidden-commands` statement at the `[edit]` hierarchy level.

## User Interface and Configuration

---

- **Monitoring, detecting, and taking action on degraded physical 10-Gigabit, 40-Gigabit, and 100-Gigabit Ethernet links to minimize packet loss (MX Series routers with MPCE, MPC3, and MPC4E)**—Starting with Junos OS Release 15.1F2, you can monitor physical link degradation (indicated by bit error rate (BER) threshold levels) on Ethernet interfaces, and take corrective actions if the BER threshold value drops to a value in the range of  $10^{-13}$  to  $10^{-5}$ .

Layer 2 and Layer 3 protocols support the monitoring of physical link degradation. An Ethernet link also supports monitoring of physical link degradation through the Link Fault Signaling (LFS) protocol. However, for both of these monitoring mechanisms, the BER threshold value range of  $10^{-13}$  to  $10^{-5}$  is very low. Because of the low BER threshold value, the physical link degradation goes undetected, causing disruption and packet loss on an Ethernet link.

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

- To monitor physical link degrade on Ethernet interfaces, configure the **link-degrade-monitor** statement.
- To configure the BER threshold value at which the corrective action must be triggered on or cleared from an interface, use the **link-degrade-monitor thresholds (set value | clear value)** statement. The *value* is the BER threshold value in scientific notation. You can configure this value in the 1E-*n* format, where 1 is the mantissa (remains constant) and *n* is the exponent. For example, a threshold value of 1E-3 refers to the BER threshold value of  $1 \times 10^{-3}$ .

The supported exponent range is 1 through 16, and the default value is 7 for the **set** configuration and 12 for the **clear** configuration.

- To configure the link degrade interval value, use the **link-degrade-monitor thresholds interval value** statement. The configured interval value determines the number of consecutive link degrade events that are considered before any corrective action is taken. The supported value range for the interval is 1 through 256, and the default interval is 10.
- To configure link degrade warning thresholds, use the **link-degrade-monitor thresholds (warning-set value | warning-clear value)** statement. The *value* is again specified in the 1E-*n* format, and the supported value range for *n* is 1 through 16. With this configuration, every time the BER threshold value is reached, a system message is logged to indicate that link degradation has occurred (**warning-set**) or link degradation has been cleared (**warning-clear**) on an interface.
- To configure the link degrade action that is taken when the configured BER threshold level is reached, use the **link-degrade action media-based** statement. A **media-based** action or fast reroute brings down the physical interface at the local end of the interface, and stops BER monitoring on the interface (although link fail is active at the local end, and recovery fail is active on the remote end of the degraded link) until an autorecovery mechanism is triggered.
- To configure the link degrade recovery options, use the **link-degrade recovery (auto interval value | manual)** statement. The recovery mechanism triggers the recovery of a degraded link.

The **auto** recovery option is used with the **media-based** action when there are no Layer 2 or Layer 3 protocols configured on the interface. With the **auto** recovery option, you must configure the **interval value** in seconds. The system then triggers the autorecovery mechanism on a degraded link. The default interval is 1800 seconds.

The **manual** recovery option is configured with the **media-based** action configuration when Layer 2 and Layer 3 protocols are configured on an interface. To trigger manual recovery, use the **request interface link-degrade-recover interface-name** statement.



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

## VPNs

- **Flow-aware transport pseudowire for BGP L2VPN and BGP VPLS (MX Series and T Series)**— Starting with Junos OS Release 15.1F2, the flow-aware transport (FAT) label that is supported for BGP-signaled pseudowires such as L2VPN and VPLS is configured only on the label edge routers (LERs). This causes the transit routers or label-switching routers (LSRs) to perform load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. The FAT flow label can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for VPWS and VPLS pseudowires.

### Related Documentation

- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 18](#)
- [Documentation Updates on page 28](#)
- [Migration, Upgrade, and Downgrade Instructions on page 28](#)
- [Product Compatibility on page 38](#)

## 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 15.1F2 for the M Series, MX Series, and T Series.

- [High Availability and Resiliency on page 9](#)

### High Availability and Resiliency

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

### Related Documentation

- [New and Changed Features on page 3](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 18](#)
- [Documentation Updates on page 28](#)

- [Migration, Upgrade, and Downgrade Instructions on page 28](#)
- [Product Compatibility on page 38](#)

## Known Behavior

There are no changes in known behavior in Junos OS Release 15.1F2 for the M Series, MX Series, and T Series.

### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 18](#)
- [Documentation Updates on page 28](#)
- [Migration, Upgrade, and Downgrade Instructions on page 28](#)
- [Product Compatibility on page 38](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1F2 for the M Series, MX Series and T Series.

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

- [Class of Service \(CoS\) on page 11](#)
- [Forwarding and Sampling on page 11](#)
- [General Routing on page 11](#)
- [Interfaces and Chassis on page 13](#)
- [Layer 2 Features on page 14](#)
- [MPLS on page 14](#)
- [Network Management and Monitoring on page 15](#)
- [Platform and Infrastructure on page 15](#)
- [Routing Protocols on page 16](#)
- [Services Applications on page 16](#)
- [Subscriber Access Management on page 16](#)
- [User Interface and Configuration on page 17](#)
- [VPNs on page 17](#)

---

### Class of Service (CoS)

---

- On MX Series platform, when class-of-service (CoS) adjustment control profiles and "overhead-accounting" are configured, if the ANCP adjust comes before the logical interface (IFL) adding message and the IFL is in "UP" state when added (for example, it may occur when carrying scaling subscribers, for instance, 8K subscribers), for some of the subscribers, the local shaping rate from dynamic profile for the subscriber logical interface may not be overridden by shaping-rate of ANCP. [PR1098006](#)

---

### Forwarding and Sampling

---

- On MX Series platform, when deleting firewall filter and the routing instance it is attached to, in some race conditions, the filter might not be deleted and remains in resolved state indefinitely. [PR937258](#)
- This defect is seen only when a existing child link from an AE is moved to a newly created AE, simultaneously from both-ends. The new AE is listed as child link in the existing AE in 'show interface ae<>.0 extensive' CLI. [PR965872](#)

---

### General Routing

---

- Periodic "show subscribers" CLI requests during the GRES recovery (on a scaled system) might lead to spawning of too many subinfo processes. As a side effect, CoA requests might not be serviced while system is kept busy by subinfo processes as authd might take long time to be recovered (it was observed that authd is not recovered after 1+ hours). [PR915677](#)
- When BCM0 interface goes down, the Routing Engine should switch over on M320. [PR949517](#)
- DHCPv6 advertise is sent with source MAC all zeroes if the subscriber is terminated on non-default routing instance. For subscribers on default instance there is no such issue observed. [PR972603](#)
- A PE device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE. The IGP instance running in the VRF on the PE may be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE. [PR977945](#)
- In point-to-point (P2P) SONET/SDH interface environment, there is a destination route with this interface as next-hop. When this interface is disabled, the destination route is still kept in the forwarding table and might cause ping fails with "Can't assign requested address" error. [PR984623](#)
- By default, the syslog utility exports 800,000 logs per second to a remote syslog server. You can modify the number of syslogs to be sent by including the message-rate-limit statement at the [edit interfaces interface-name services-options syslog] hierarchy level to suit your deployment needs. The rate at which syslog messages can be sent to the Routing Engine is 10,000 logs per second. [PR1001201](#)
- On MX Series platform with IPsec, NAT or RTM service configured on MS-MPC or MS-MIC inserted in an XM chip based MPC, if the packet undergoes fragmentation

over the multiservice (ms) interface and the last fragment is very small (the last fragment packet length is less than 120 bytes), the MPC may crash and reboot due to the packet length calculation defect. As a workaround, we can disable fragmentation over the ms interface. [PR1025824](#)

- On MX Series with MPCs and MICs, in Connection-less Network Service (CLNS) L3VPN over Ethernet scenario, because the IS-IS route may be incorrectly treated as ARP-type route (32 bits routes with the MAC and outgoing interface specified in the nexthop), the routing protocol process (rpd) of Provider Edge (PE) device may reject to add the route. [PR1041251](#)
- When enabling pseudowire subscribers the "show subscribers extensive" command does not display CoS policies applied to the subscriber interface. [PR1060036](#)
- Output MTU counter shows incorrect data in the show pfe statistics traffic command output. [PR1061111](#)
- For MX-VC platform, performing unified ISSU in scaled subscribers environment might cause all VC members to get restarted unexpectedly. [PR1070542](#)
- On MX Series routers with MPC based line cards in a setup involving Packet Forwarding Engine fast reroute (FRR) applications, when BFD session flaps the next-hop program in the Packet Forwarding Engine may get corrupted. It may lead to incorrect selection of next-hop or traffic blackhole. [PR1071028](#)
- Traffic throughput test between MPC1/1E/2/2E card and MPC2E/3E NG card, flowing from MPC1/1E/2/2E card to MPC2E/3E NG card is less than from MPC2E/3E NG card to MPC1/1E/2/2E card. [PR1076009](#)
- In subscriber management environment, the PPP daemon (jpppd) might crash repeatedly due to a memory double-free issue. [PR1079511](#)
- 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](#)
- TCP messages do not have their MSS adjusted by the MS-MPC (SPARKS) card, if they do not belong to an established session. [PR1084653](#)
- After upgrading to Release 14.2R3.8, loopback ISO family address may be stuck in KRT queue. [PR1097778](#)
- Neighbor deletion and re-creation on a session when back pressure is applied on all the active neighbor sessions could cause the new session to be in stuck state. This PR addresses this problem by ensuring that the neighbor session is established first. [PR1098549](#)
- If a loopback logical interface has greater than 400 addresses configured, we will have memory leak in bbe-smgd. This bug exists in 14.1X51-D50. [PR1101652](#)
- In VRRP for IPv6 environment, when feature "forwarding IPv6 Solicited Router Advertisement as unicast" is enabled (by enabling configuration statement "solicit-router-advertisement-unicast"), if the configuration statement "virtual-router-only" is also configured on the interface, the interface (if the groups are

in the master state) may not respond with router-advertisement (RA) messages after receiving the router-solicit messages (RS). [PR1103113](#)

- If fpc offline configuration statement is configured after the presence of Non-recoverable faults, then offline action will not be performed. [PR1103185](#)
- 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 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](#)

## Interfaces and Chassis

- Packet Forwarding Engine continues to forward traffic to DHCP client on a demux interface when ae0 interface is down. In this scenario the AE interface bundle has five members and configured with minimum link value of 4. When two members are down, the ae0 interface also goes down, but Packet Forwarding Engine continues to forward traffic on other members for the demux interface. [PR836846](#)
- Chap Local-name defaults to 8 characters. Should be 32. [PR996760](#)
- In rare case, when a child link flaps within an aggregate bundle which happens twice within a short period of time (that is, if the child interface comes up within a short period of time after it has gone down), there is a probability that a race condition might happen. The result is to have child next hop within aggregate next hop to be in "Replaced" state on the FPC, leading to traffic blackholing. [PR1032931](#)
- In some configurations `agg_pfe_get_fwd_options`, log message is generated at the excessive rate. This log message can be helpful during troubleshooting, but it is not needed during normal operation. Though it is not impacting service, it may increase the load of the system and it was decided to cover this message under traceoptions in order to optimize system performance. [PR1047564](#)
- 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 ppeid 0xe-0/2/0: get tlv ppeid 0xe-0/3/0: get tlv ppeid 0xe-1/2/0: get tlv ppeid 0xe-1/3/0: get tlv ppeid 0xe-2/0/0: get tlv ppeid 0xe-2/1/0: get tlv ppeid 0xe-2/2/0: get tlv ppeid 0xe-2/3/0: get tlv ppeid 0xe-5/1/0: get tlv ppeid 0xe-5/1/1: get tlv ppeid 0xe-5/1/2: get tlv ppeid 0xe-5/1/3: get tlv ppeid 0xe-5/1/4: get tlv ppeid 0xe-5/1/5: get tlv ppeid 0xe-5/1/6: get tlv ppeid 0xe-5/1/7: get tlv ppeid 0xe-5/1/8: get tlv ppeid 0xe-5/1/9: get tlv ppeid 0. [PR1057594](#)
- show interface extensive will show these additional counters... <<<< Preclassifier statistics: Traffic Class Received Packets Transmitted Packets Dropped Packets real-time 0 0 3 network-control 0 0 3 best-effort 0 0 3 <<<< [PR1060070](#)
- Link Up/Down SNMP traps for AE member links might not be generated, but the SNMP traps for the AE bundle works well. [PR1067011](#)

- Deactivating/activating logical interfaces may cause BGP session flap when BGP is using VRRP VIP as source address. This is caused by a timing issue between DCD and VRRP overlay file. When DCD reads the overlay file, it is not the updated one or yet to be updated. This results in error and DCD stops parsing VRRP overlay file. [PR1089576](#)
- On MX Series platform, "cfp\_lh\_update\_lsec\_pm\_var received" messages are periodically logged with Warning level. The severity of this message has been revised. [PR1089592](#)
- In the dual Routing Engines scenario with GRES and ae0 interfaces configuration, if GRES is disabled on system, the backup Routing Engine should remove the ae0 bundle. However, it does not go clean and ae0 remains in backup Routing Engine. After switching, Routing Engine mastership to make other Routing Engine as master, the new master Routing Engine (which was backup earlier) continues to use invalid MAC address "00:00:00:00:00:00". [PR1089946](#)

---

## Layer 2 Features

- When "input-vlan-map" with "push" operation is enabled for dual-tagged interfaces in "enhanced-ip" mode, there is a probability that the broadcast, unknown unicast, and multicast (BUM) traffic may be blackholed on some of the child interfaces of the egress Aggregated Ethernet (AE) interfaces. [PR1078617](#)
- V44 defines the next generation of Juniper Networks Fabric using MX as the aggregation device (AD) and EX4300/QFX5100's as the Satellite Devices (SD). When V44 port extension is in use, after switching from Master to Backup Routing Engine, the pppman daemon on the SDs may crash and not be restarted. It results in the aggregated Ethernet (ae) bundle over v44 extended ports does not come up. [PR1101266](#)

---

## 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 gets isolated, unexpected packet drops will be observed. [PR1030815](#)
- In Resource Reservation Protocol (RSVP) environment, if CoS-based forwarding (CBF) for per LSP (that filter out traffic not related to that LSP) is configured, and either the feature fast-reroute or link-protection is used on the device, when the primary link is down (for example, turning off the laser of the link), due to some next hops of the traffic may be deleted or reassigned to different class of traffic, the RSVP local repair may fail to process more than 200 LSPs at one time. The traffic may get dropped by the filter on the device before the new next hop is installed. In this situation, the feature (fast reroute or link protection) may take longer time (for example, 1.5 seconds) to function and the traffic loss might be seen at the meantime. In addition, the issue may not be seen if the CBF for per LSP is not configured on the device. [PR1048109](#)

## 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 MX Series-only filter applied. > Jul 8 15:52:09 galway-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) > Jul 8 15:52:09 galway-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) > Jul 8 15:52:09 galway-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) > Jul 8 15:52:09 galway-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). [PR988566](#)

## Platform and Infrastructure

- MX Series-based line card might crash when trying to install the composite next-hop used for the next-hop-group configuration related to port mirroring of traffic over IRB to an LSI attached to VPLS instance for a remote host. [PR1029070](#)
- On MX Series-based platform, when using inline Two-Way Active Measurement Protocol (TWAMP) server (the server address is the inline service interface address), because the TWAMP server may incorrectly calculate the packet checksum, the packet may get dropped on the TWAMP client. [PR1042132](#)
- When using the "ping detail" command, the interface number is provided on the output instead of the interface name. [PR1078300](#)
- 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](#)
- If with both MPC/MSDPC and other type of DPCs equipped, for local switching at mesh group level, split horizon on PW interfaces won't work and this would cause packets to loop back to same PW interface. [PR1084130](#)

- IPv6 packets with non-UDP and non-TCP payload belonging to the same flow might get re-ordered when being forwarded by MX Series MPC Packet Forwarding Engine. [PR1098776](#)
- Due to a software defect found in 13.3R7.3 and 14.1R5.4 inclusively, Juniper Networks strongly discourage the use of Junos software version 13.3R7.3 on routers with MQ-based MPC. This includes MX-Series with MPC1, MPC2; all mid-range MX-Series; and some of EX9200 line cards. [PR1108826](#)

---

### Routing Protocols

- It is necessary that the MSDP peer local-address matches the PIM RP address on routers that are RP. MSDP RPF check might fail in rare cases when both these addresses are not equal. [PR35806](#)
- Continuous soft core-file may be observed due to bgp-path-selection code. RPD forks a child and the child asserts to produce a core-file. The problem is with route-ordering and it is auto-corrected after collecting the soft-assert-corefile, without any impact to traffic/service. [PR815146](#)
- Support for the Pragmatic General Multicast protocol (daemon pgmd) is being phased out from Junos OS. In Junos OS Release 14.2 the CLI is now hidden (although the component is still there and configurable). In Junos OS Release 15.1 the code and its corresponding CLI are removed. [PR936723](#)
- 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](#)

---

### Services Applications

- In the NAT environment, the jnxNatSrcPoolName OID is not implemented in jnxSrcNatStatsTable. [PR1039112](#)
- In some cases after unified ISSU upgrade/GRES switch/jl2tpd restart if the subscriber is terminated during the unified ISSU/GRES/restart process jl2tpd may core. This PR addresses/fixes this issue. [PR1109447](#)
- With scaling Layer 2 Tunneling Protocol (L2TP) sessions (for example, 128k sessions), when executing L2TP "show" command in one terminal and "clear" command in another terminal simultaneously, pressing Ctrl-C or closing the terminal on one terminal might cause the jl2tpd process crash. [PR1063207](#)

---

### Subscriber Access Management

- When the MX router acting as the Policy and Charging Enforcement Function (PCEF) uses Gx-Plus to request service provisioning from the Policy Control and Charging Rules Function (PCRF), the authentication service process (authd) might crash during the subscribers logout. [PR1034287](#)
- This issue was introduced as part of another fix. Please contact JTAC for the recommended release for your deployment. [PR1049955](#)
- authd core at /src/junos/usr.sbin/authd/authd\_ipc.c:1145 with JSRC. [PR1094674](#)



## 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](#)

## VPNs

- When you modify the frame-relay-tcc statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level of a Layer 2 VPN, the connection for the second logical interface might not come up. As a workaround, restart the chassis process (chassisd) or reboot the router. [PR32763](#)
- It is planned that future releases of Junos OS will modify the default BGP extended community value used for MVPN IPv4 VRF Route Import (RT-Import) to the IANA-standardized value. Thus, default behavior is expected to change such that the behavior of the configuration 'mvpn-iana-rt-import' will become the default and the 'mvpn-iana-rt-import' configuration will be deprecated. [PR890084](#)

### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Resolved Issues on page 18](#)
- [Documentation Updates on page 28](#)
- [Migration, Upgrade, and Downgrade Instructions on page 28](#)
- [Product Compatibility on page 38](#)

## Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases.

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

- [Resolved Issues: 15.1F2 on page 18](#)

---

### Resolved Issues: 15.1F2

- [Class of Service \(CoS\) on page 18](#)
- [General Routing on page 19](#)
- [Infrastructure on page 21](#)
- [Interfaces and Chassis on page 21](#)
- [Layer 2 Features on page 23](#)
- [MPLS on page 23](#)
- [Network Management and Monitoring on page 24](#)
- [Platform and Infrastructure on page 24](#)
- [Routing Protocols on page 26](#)
- [Services Applications on page 26](#)
- [Software Installation and Upgrade on page 27](#)
- [User Interface and Configuration on page 27](#)
- [VPNs on page 27](#)

#### ***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 CLI command "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 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 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 (vrpd) 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 config 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 It (logical tunnel) interface, if deactivating and activating scaled configuration a few times, in rare condition, the It 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 It 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 needs 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](#)
- 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. MX Series, T4000, and PTX Series 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. It might cause the LDP session between routers to go down. [PR1065338](#)
- 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 statement 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 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 Trio based 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 Trio-based line cards. 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](#)
- Starting from Junos OS Release 14.2R1, the CLI command "set date ntp a.b.c.d" may not be working. [PR1067107](#)
- StartTime and EndTime of the flow in inline-jflow (version 9) has future time-stamp. [PR1067107](#)
- 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 Software flow into account when calculating the flow limit. [PR1023439](#)
- On M Series, MX Series, T Series routers with Multiservices 100, Multiservices 400, or Multiservices 500 PICs with "dump-on-flow-control" configured, if prolonged flow control failure, the coredump file might generate failure. [PR1039340](#)

- On MX Series routers that are acting as LNS to provide tunnel endpoints, it is observed that the service-interfaces are not usable if a MIC corresponding to them is not physically installed on the FPC. If only those service interfaces that belong to the removed PIC are added to service-device-pool, this results in no LNS subscribers being able to log in. Note that once the MIC is inserted into the FPC, the features could be used. [PR1063024](#)
- When configuring RADIUS authentication for Layer 2 Tunneling Protocol (L2TP), the RADIUS server cannot be recognized because the source address is not being read correctly. As a result, the L2TP session cannot be established. [PR1064817](#)
- The trigger for the crash is when the MS-DPCs Service PIC is in a low memory zone and it receives two SYN messages from the the same client IP within a very short time gap in between the two SYNs. So this race condition is tied to running out of memory, failing to allocating a timer for a conversation, and having rapid SYNs on a TCP connection where the second TCP SYN is matched on flow which is being deleted due to a failed timer allocation for that. This scenario is very difficult to hit and should not be seen in production often. [PR1069006](#)
- Service PIC daemon (spd) might crash with core-dumps due to CGNAT pool's snmp-trap-thresholds configuration. [PR1070370](#)
- Earlier output from "show service l2tp tunnel" will not display tunnels with no sessions. This behavior have been changed, now empty tunnels are also displayed in this command. [PR1071923](#)

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

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

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

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

### ***VPNs***

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

### **Related Documentation**

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)

- [Known Issues on page 10](#)
- [Documentation Updates on page 28](#)
- [Migration, Upgrade, and Downgrade Instructions on page 28](#)
- [Product Compatibility on page 38](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 15.1F2 documentation for the M Series, MX Series, and T Series.

### Related Documentation

- [New and Changed Features on page 3](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 28](#)
- [Product Compatibility on page 38](#)

## 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 M Series, 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:** In Junos OS Release 15.1, Junos OS (FreeBSD 10.x) is not available to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to use the existing Junos OS (FreeBSD 6.1).



**CAUTION:** This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for M Series, MX Series, and T Series routers. We recommend that you not use 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
M7i, M10i, M120, M320	YES	NO
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 15.1F2 on page 29](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) on page 31](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 32](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 34](#)
- [Upgrading a Router with Redundant Routing Engines on page 34](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 34](#)
- [Upgrading the Software for a Routing Matrix on page 36](#)
- [Upgrading Using Unified ISSU on page 37](#)
- [Downgrading from Release 15.1 on page 37](#)

### **Basic Procedure for Upgrading to Release 15.1F2**

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.



**NOTE:** This section does not apply to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to refer to the next section.

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.

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-15.1F2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-15.1F2.9-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)

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 15.1F2 **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 M Series routers, all T Series routers, MX80, and MX104.



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

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-15.1F2.9-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-15.1F2.9-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 15.1F2 **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 Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1

---

In releases prior to Junos OS Release 10.1, the draft-rosen multicast VPN feature implements the unicast lo0.x address configured within that instance as the source address used to establish PIM neighbors and create the multicast tunnel. In this mode, the multicast VPN loopback address is used for reverse path forwarding (RPF) route resolution to create the reverse path tree (RPT), or multicast tunnel. The multicast VPN loopback address is also used as the source address in outgoing PIM control messages.

In Junos OS Release 10.1 and later, you can use the router's main instance loopback (lo0.0) address (rather than the multicast VPN loopback address) to establish the PIM state for the multicast VPN. We strongly recommend that you perform the following procedure when upgrading to Junos OS Release 10.1 if your draft-rosen multicast VPN network includes both Juniper Network routers and other vendors' routers functioning as provider edge (PE) routers. Doing so preserves multicast VPN connectivity throughout the upgrade process.

Because Junos OS Release 10.1 supports using the router's main instance loopback (lo0.0) address, it is no longer necessary for the multicast VPN loopback address to match the main instance loopback address lo0.0 to maintain interoperability.



**NOTE:** You might want to maintain a multicast VPN instance lo0.x address to use for protocol peering (such as IBGP sessions), or as a stable router identifier, or to support the PIM bootstrap server function within the VPN instance.

Complete the following steps when upgrading routers in your draft-rosen multicast VPN network to Junos OS Release 10.1 if you want to configure the routers's main instance loopback address for draft-rosen multicast VPN:

1. Upgrade all M7i and M10i routers to Junos OS Release 10.1 before you configure the loopback address for draft-rosen Multicast VPN.



**NOTE:** Do not configure the new feature until all the M7i and M10i routers in the network have been upgraded to Junos OS Release 10.1.

2. After you have upgraded all routers, configure each router's main instance loopback address as the source address for multicast interfaces.

Include the **default-vpn-source interface-name loopback-interface-name** statement at the **[edit protocols pim]** hierarchy level.

3. After you have configured the router's main loopback address on each PE router, delete the multicast VPN loopback address (lo0.x) from all routers.

We also recommend that you remove the multicast VPN loopback address from all PE routers from other vendors. In Junos OS releases prior to 10.1, to ensure interoperability with other vendors' routers in a draft-rosen multicast VPN network, you had to perform additional configuration. Remove that configuration from both the Juniper Networks routers and the other vendors' routers. This configuration should be on Juniper Networks routers and on the other vendors' routers where you configured the lo0.mvpn address in each VRF instance as the same address as the main loopback (lo0.0) address.

This configuration is not required when you upgrade to Junos OS Release 10.1 and use the main loopback address as the source address for multicast interfaces.



**NOTE:** To maintain a loopback address for a specific instance, configure a loopback address value that does not match the main instance address (lo0.0).

For more information about configuring the draft-rosen Multicast VPN feature, see the [Multicast Protocols Feature Guide for Routing Devices](#).

### 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 M Series, MX Series, and T Series routers. We recommend that you not use 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 3](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 18](#)
- [Documentation Updates on page 28](#)
- [Product Compatibility on page 38](#)

## Product Compatibility

- [Hardware Compatibility on page 38](#)

### 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 M Series, MX Series, and T 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 3](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 10](#)
- [Resolved Issues on page 18](#)
- [Documentation Updates on page 28](#)
- [Migration, Upgrade, and Downgrade Instructions on page 28](#)

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

---

These release notes accompany Junos OS Release 15.1F2 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 39](#)
- [Changes in Behavior and Syntax on page 40](#)
- [Known Behavior on page 41](#)
- [Known Issues on page 41](#)
- [Resolved Issues on page 42](#)
- [Documentation Updates on page 44](#)
- [Migration, Upgrade, and Downgrade Instructions on page 44](#)
- [Product Compatibility on page 47](#)

### New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1F2 for the PTX Series.

- [Class of Service \(CoS\) on page 39](#)
- [Multicast on page 40](#)
- [VPNs on page 40](#)

#### Class of Service (CoS)

---

- **Change in scaling number for rewrite rules on PTX Series Routers**—Starting with Release 15.1F2, on PTX Series routers, the scaling number for a rewrite rule is reduced by one when the default EXP rewrite is used. This change in scaling number is introduced to:
  - Support all possible combinations of EXP rewrite rules.

- Fix the issue of incorrect modification of EXP bits of the inner label by the default MPLS EXP rewrite rule during the label pop operation.

## Multicast

---

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

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

## VPNs

---

- **Flow-aware transport pseudowire for BGP L2VPN and BGP VPLS (PTX Series)**—Starting with Junos OS Release 15.1F2, the flow-aware transport (FAT) label that is supported for BGP-signaled pseudowires such as L2VPN and VPLS is configured only on the label edge routers (LERs). This causes the transit routers or label-switching routers (LSRs) to perform load balancing of MPLS packets across equal-cost multipath (ECMP) paths or link aggregation groups (LAGs) without the need for deep packet inspection of the payload. The FAT flow label can be used for LDP-signaled forwarding equivalence class (FEC 128 and FEC 129) pseudowires for VPWS and VPLS pseudowires.

### Related Documentation

- [Changes in Behavior and Syntax on page 40](#)
- [Known Behavior on page 41](#)
- [Known Issues on page 41](#)
- [Resolved Issues on page 42](#)
- [Documentation Updates on page 44](#)
- [Migration, Upgrade, and Downgrade Instructions on page 44](#)
- [Product Compatibility on page 47](#)

## Changes in Behavior and Syntax

There are no changes in default behavior and syntax in Junos OS Release 15.1F2 for the PTX Series.

### Related Documentation

- [New and Changed Features on page 39](#)
- [Known Behavior on page 41](#)
- [Known Issues on page 41](#)
- [Resolved Issues on page 42](#)
- [Documentation Updates on page 44](#)
- [Migration, Upgrade, and Downgrade Instructions on page 44](#)



- [Product Compatibility on page 47](#)

## Known Behavior

There are no changes in known behavior in Junos OS Release 15.1F2 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 39](#)
- [Changes in Behavior and Syntax on page 40](#)
- [Known Issues on page 41](#)
- [Documentation Updates on page 44](#)
- [Resolved Issues on page 42](#)
- [Migration, Upgrade, and Downgrade Instructions on page 44](#)
- [Product Compatibility on page 47](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1F2 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 41](#)
- [Interfaces and Chassis on page 41](#)
- [MPLS on page 42](#)
- [Routing Protocols on page 42](#)

### General Routing

- The PTX Series does not support the queuing PICs. However, 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 qpik 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](#)

### Interfaces and Chassis

- In rare case, when a child link flaps within an aggregate bundle which happens twice within a short period of time (that is, if the child interface comes up within a short period of time after it has gone down), there is a probability that a race condition might happen.

The result is to have child NH within aggregate NH to be in "Replaced" state on the FPC, leading to traffic blackholing. [PR1032931](#)

- On PTX Series routers, optical threshold value is shown incorrectly for the interfaces in the PIC P1-PTX-2-100G-WDM. [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](#)

## 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](#)

## 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. The severity of the log has been reduced to INFO. [PR1065055](#)
- In IS-IS environment, MPLS LSPs are established, and when IS-IS traceoptions flag "general" is activated, the LSP convergence time is increased. [PR1090752](#)

### Related Documentation

- [New and Changed Features on page 39](#)
- [Changes in Behavior and Syntax on page 40](#)
- [Known Behavior on page 41](#)
- [Resolved Issues on page 42](#)
- [Documentation Updates on page 44](#)
- [Migration, Upgrade, and Downgrade Instructions on page 44](#)
- [Product Compatibility on page 47](#)

## Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases. The identifier following the description is the tracking number in the Juniper Networks Problem Report (PR) tracking system.

- [Resolved Issues: 15.1F2 on page 42](#)

### Resolved Issues: 15.1F2

---

- [Forwarding and Sampling on page 43](#)
- [General Routing on page 43](#)
- [Interfaces and Chassis on page 43](#)
- [MPLS on page 43](#)

### **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. MX Series, T4000, and PTX Series platform 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 39](#)
  - [Changes in Behavior and Syntax on page 40](#)
  - [Known Issues on page 41](#)
  - [Documentation Updates on page 44](#)
  - [Known Behavior on page 41](#)
  - [Migration, Upgrade, and Downgrade Instructions on page 44](#)
  - [Product Compatibility on page 47](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 15.1F2 documentation for the PTX Series.

- Related Documentation**
- [New and Changed Features on page 39](#)
  - [Changes in Behavior and Syntax on page 40](#)
  - [Known Behavior on page 41](#)
  - [Known Issues on page 41](#)
  - [Resolved Issues on page 42](#)
  - [Migration, Upgrade, and Downgrade Instructions on page 44](#)
  - [Product Compatibility on page 47](#)

## 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 44](#)
- [Upgrading a Router with Redundant Routing Engines on page 45](#)
- [Basic Procedure for Upgrading to Release 15.1F2 on page 45](#)

### 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 15.1F2

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 15.1F2 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 15.1F2 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-15.1F21-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-15.1
F21-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 15.1F2 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 39](#)
- [Changes in Behavior and Syntax on page 40](#)
- [Known Behavior on page 41](#)
- [Known Issues on page 41](#)
- [Resolved Issues on page 42](#)
- [Documentation Updates on page 44](#)
- [Product Compatibility on page 47](#)

## Product Compatibility

- [Hardware Compatibility on page 48](#)

## 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 39](#)
- [Changes in Behavior and Syntax on page 40](#)
- [Known Behavior on page 41](#)
- [Known Issues on page 41](#)
- [Resolved Issues on page 42](#)
- [Documentation Updates on page 44](#)
- [Migration, Upgrade, and Downgrade Instructions on page 44](#)



## 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 provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at <http://www.juniper.net/techpubs/index.html>, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at <http://www.juniper.net/techpubs/feedback/>.
- E-mail—Send your comments to [techpubs-comments@juniper.net](mailto:techpubs-comments@juniper.net). Include the document or topic name, URL or page number, and software 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

7 January 2016—Revision 2, Junos OS Release 15.1F2— M Series, MX Series, PTX Series, and T Series.

21 August 2015—Revision 1, Junos OS Release 15.1F2— M Series, 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.