

# Release Notes:Junos<sup>®</sup> OS Release 15.1R2 for the EX Series, M Series, MX Series, PTX Series, and T Series

18 February 2016

## Contents

Introduction .....	6
Junos OS Release Notes for EX Series Switches .....	6
New and Changed Features .....	6
Interfaces and Chassis .....	7
Junos OS XML API and Scripting .....	8
Management .....	8
MPLS .....	9
Port Security .....	9
Software Installation and Upgrade .....	10
Spanning-Tree Protocols .....	10
Changes in Behavior and Syntax .....	10
Dynamic Host Configuration Protocol .....	11
Known Behavior .....	11
Authentication and Access Control .....	11
J-Web .....	12
Port Security .....	12
Spanning-Tree Protocols .....	12
Virtual Chassis .....	12
Known Issues .....	13
Dynamic Host Configuration Protocol .....	13
Infrastructure .....	13
Interfaces and Chassis .....	14
J-Web .....	14
Software Installation and Upgrade .....	14
Resolved Issues .....	15
Resolved Issues: Release 15.1R2 .....	15
Documentation Updates .....	18
Migration, Upgrade, and Downgrade Instructions .....	18
Upgrade and Downgrade Support Policy for Junos OS Releases .....	18

Product Compatibility . . . . .	19
Hardware Compatibility . . . . .	19
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers . . . . .	21
New and Changed Features . . . . .	21
Hardware . . . . .	22
Bridging and Learning . . . . .	23
Class of Service (CoS) . . . . .	23
High Availability (HA) and Resiliency . . . . .	24
Interfaces and Chassis . . . . .	26
IPv6 . . . . .	30
Junos OS XML API and Scripting . . . . .	30
Layer 2 Features . . . . .	30
Management . . . . .	31
MPLS . . . . .	32
Multicast . . . . .	33
Network Management and Monitoring . . . . .	34
Routing Policy and Firewall Filters . . . . .	36
Routing Protocols . . . . .	37
Services Applications . . . . .	39
Software Defined Networking . . . . .	43
Software Installation and Upgrade . . . . .	43
Subscriber Management and Services (MX Series) . . . . .	44
User Interface and Configuration . . . . .	50
VPNs . . . . .	50
Changes in Behavior and Syntax . . . . .	52
Authentication, Authorization and Accounting . . . . .	53
Class of Service (CoS) . . . . .	53
General Routing . . . . .	53
High Availability (HA) and Resiliency . . . . .	53
Junos XML API and Scripting . . . . .	55
Layer 2 VPNs . . . . .	55
MPLS . . . . .	55
Multicast . . . . .	55
Network Management and Monitoring . . . . .	55
Routing Policy and Firewall Filters . . . . .	56
Routing Protocols . . . . .	56
Security . . . . .	59
Services Applications . . . . .	59
Subscriber Management and Services (MX Series) . . . . .	60
System Logging . . . . .	65
System Management . . . . .	66
User Interface and Configuration . . . . .	66
VPNs . . . . .	66
Known Behavior . . . . .	67
Hardware . . . . .	67
MPLS . . . . .	68
Subscriber Management and Services (MX Series) . . . . .	68
System Logging . . . . .	68

Known Issues	68
Forwarding and Sampling	69
General Routing	69
Infrastructure	71
Interfaces and Chassis	71
Layer 2 Features	72
MPLS	72
Network Management and Monitoring	72
Platform and Infrastructure	73
Routing Protocols	74
Services Applications	74
Software Installation and Upgrade	75
Subscriber Access Management	75
User Interface and Configuration	75
VPNs	75
Resolved Issues	76
Class of Service (CoS)	76
Forwarding and Sampling	77
General Routing	80
High Availability (HA) and Resiliency	84
Interfaces and Chassis	84
Layer 2 Features	88
MPLS	89
Network Management and Monitoring	89
Platform and Infrastructure	90
Routing Policy and Firewall Filters	94
Routing Protocols	94
Services Applications	96
Software Installation and Upgrade	97
Subscriber Access Management	97
User Interface and Configuration	98
VPNs	98
Documentation Updates	99
Adaptive Services Interfaces Feature Guide for Routing Devices	99
Broadband Subscriber Sessions Feature Guide	100
Broadband Subscriber VLANs and Interfaces Feature Guide	100
High Availability Feature Guide	100
IPv6 Neighbor Discovery Feature Guide for Routing Devices	101
Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices	101
MPLS Applications Feature Guide for Routing Devices	102
Overview for Routing Devices	103
Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices	103
Security Services Administration Guide for Routing Devices	103
User Access and Authentication Guide for Routing Devices	103
VPNs Library for Routing Devices	103

Migration, Upgrade, and Downgrade Instructions . . . . .	104
Basic Procedure for Upgrading to Release 15.1 . . . . .	105
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x) . . . . .	106
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1) . . . . .	107
Upgrade and Downgrade Support Policy for Junos OS Releases . . . . .	109
Upgrading a Router with Redundant Routing Engines . . . . .	109
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 . . . . .	109
Upgrading the Software for a Routing Matrix . . . . .	111
Upgrading Using Unified ISSU . . . . .	112
Downgrading from Release 15.1 . . . . .	112
Product Compatibility . . . . .	113
Hardware Compatibility . . . . .	113
Junos OS Release Notes for PTX Series Packet Transport Routers . . . . .	114
New and Changed Features . . . . .	114
High Availability and Resiliency (HA) . . . . .	115
Interfaces and Chassis . . . . .	115
IPv6 . . . . .	116
Junos OS XML API and Scripting . . . . .	116
Management . . . . .	117
MPLS . . . . .	118
Routing Protocols . . . . .	118
User Interface and Configuration . . . . .	119
VPNs . . . . .	120
Changes in Behavior and Syntax . . . . .	120
High Availability (HA) and Resiliency . . . . .	120
Junos OS XML API and Scripting . . . . .	121
Routing Protocols . . . . .	121
User Interface and Configuration . . . . .	121
Known Behavior . . . . .	122
System Logging . . . . .	122
Known Issues . . . . .	122
General Routing . . . . .	122
Interfaces and Chassis . . . . .	123
Network Management and Monitoring . . . . .	123
Routing Protocols . . . . .	123
Software Installation and Upgrade . . . . .	124
Resolved Issues . . . . .	124
Forwarding and Sampling . . . . .	125
General Routing . . . . .	125
Interfaces and Chassis . . . . .	126
MPLS . . . . .	127
Network Management and Monitoring . . . . .	127
Routing Protocols . . . . .	127
Documentation Updates . . . . .	127
High Availability Feature Guide . . . . .	127
IPv6 Neighbor Discovery Feature Guide . . . . .	128

---

Migration, Upgrade, and Downgrade Instructions . . . . .	128
Upgrading Using Unified ISSU . . . . .	128
Upgrading a Router with Redundant Routing Engines . . . . .	129
Basic Procedure for Upgrading to Release 15.1 . . . . .	129
Product Compatibility . . . . .	131
Hardware Compatibility . . . . .	132
Third-Party Components . . . . .	133
Finding More Information . . . . .	133
Documentation Feedback . . . . .	133
Requesting Technical Support . . . . .	134
Self-Help Online Tools and Resources . . . . .	134
Opening a Case with JTAC . . . . .	134
Revision History . . . . .	135

## 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 15.1R2 for the EX Series, 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 EX Series Switches

---

These release notes accompany Junos OS Release 15.1R2 for the EX 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/>.

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 10](#)
- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R2 for the EX Series.



**NOTE:** The following EX Series platforms are supported in Release 15.1R2: EX2200, EX3300, EX4200, EX4500, EX4550, EX6200, EX8200, and EX9200.



**NOTE:** A new J-Web distribution model was introduced in Junos OS Release 14.1X53-D10, and that same model is supported in Release 15.1R1. The model provides two packages:

- Platform package—Installed as part of Junos OS; provides basic functionalities of J-Web.
- Application package—Optionally installable package; provides complete functionalities of J-Web.

The J-Web Platform package is included in the EX2200, EX3300, EX4200, EX4500, EX4550, and EX6210 Junos OS Release 15.1R1 install images.

For details about the J-Web distribution model, see [Release Notes: J-Web Application Package Release 15.1A1 for Juniper Networks EX Series Ethernet Switches](#).

- [Interfaces and Chassis on page 7](#)
- [Junos OS XML API and Scripting on page 8](#)
- [Management on page 8](#)
- [MPLS on page 9](#)
- [Port Security on page 9](#)
- [Software Installation and Upgrade on page 10](#)
- [Spanning-Tree Protocols on page 10](#)

## Interfaces and Chassis

- **Support for MC-LAG on logical systems (EX9200 switches)**—Starting with Junos OS Release 15.1, you can configure multichassis link aggregation (MC-LAG) interfaces on logical systems within an EX9200 switch. When you configure multichassis aggregated Ethernet interfaces on a logical system, ensure that these interfaces are added with the same multichassis aggregated Ethernet identification number and redundancy group identifier for the MC-LAG on both peers or devices that are connected by the MC-AE interfaces. Ensure that the Inter-Chassis Control Protocol (ICCP) to associate the routing or switching devices contained in a redundancy group is defined on both peers within the logical systems of the devices. Such a configuration ensures that all packets are transmitted using ICCP within the logical system network. The logical system information is added, and then removed, by the ICCP process to prevent each packet from containing the logical system details. This behavior enables multiple disjoint users to employ MC-LAG capabilities within their networks transparently and seamlessly. A unique ICCP definition for a logical system is created, thereby enabling you to wholly manage ICCP parameters on one logical system without the need for access permissions to view other logical system networks on the same device.

Configuration of MC-LAG interfaces on logical systems enables MC-LAG to be used across multiple routing tables and switch forwarding tables in active-active and active-standby modes of MC-LAG interfaces.

[See [Multichassis Link Aggregation on Logical Systems Overview](#).]

- **IPv6 support on multichassis aggregated Ethernet interfaces (EX9200 switches)**—Starting with Junos OS Release 15.1, multichassis aggregated Ethernet interfaces on EX9200 switches support IPv6 and Neighbor Discovery Protocol (NDP). IPv6 neighbor discovery is a set of ICMPv6 messages that combine IPv4 messages such as ICMP redirect, ICMP router discovery, and ARP messages.

[See [Understanding IPv6 Neighbor Discovery Protocol and MC-LAGs on EX9200 Switches](#).]

## Junos OS XML API and Scripting

---

- **Support for replacing patterns in configuration data within NETCONF and Junos XML protocol sessions (EX Series)**—Starting with Junos OS Release 15.1, you can replace variables and identifiers in the candidate configuration when you perform a `<load-configuration>` operation in a Junos XML protocol or NETCONF session. The `replace-pattern` attribute specifies the pattern to replace, the `with` attribute specifies the replacement pattern, and the optional `upto` attribute indicates the number of occurrences to replace. The scope of the replacement is determined by the placement of the attributes in the configuration data. The functionality of the attribute is identical to that of the `replace pattern` configuration mode command in the Junos OS CLI.

[See [Replacing Patterns in Configuration Data Using the NETCONF or Junos XML Protocol](#).]

## Management

---

- **Support for YANG features, including configuration hierarchy must constraints published in YANG, and a module that defines Junos OS YANG extensions (EX Series)**—Starting with Junos OS Release 15.1, the Juniper Networks `configuration` YANG module includes configuration constraints published using either the YANG `must` statement or the Junos OS YANG extension `junos:must`. Constraints that cannot be mapped directly to the YANG `must` statement, which include expressions containing special keywords or symbols such as `all`, `any`, `unique`, `$`, `__`, and wildcard characters, are published using `junos:must`.

The `junos-extension` module contains definitions for Junos OS YANG extensions, including the `must` and `must-message` keywords. The `junos-extension` module is bound to the namespace URI `http://yang.juniper.net/yang/1.1/je` and uses the prefix `junos`. You can download Juniper Networks YANG modules from the website, or you can generate the modules by using the `show system schema` operational mode command on your local device.

[See [Using Juniper Networks YANG Modules](#).]

- **Support for enforcing RFC-compliant behavior in NETCONF sessions (EX Series)**—Starting with Junos OS Release 15.1, you can require that the NETCONF server enforce certain behaviors during the NETCONF session by configuring the `rfc-compliant` statement at the `[edit system services netconf]` hierarchy level. If you configure the `rfc-compliant` statement, the NETCONF server explicitly declares the NETCONF namespace in its replies and qualifies all NETCONF tags with the `nc` prefix. Also, `<get>` and `<get-config>` operations that return no configuration data do not include an empty `<configuration>` element in RPC replies.



[See [Configuring RFC-Compliant NETCONF Sessions.](#)]

## MPLS

---

- **New command to display the MPLS label availability in RPD (EX Series)**—Starting with Junos OS Release 15.1, a new show command, **show mpls label usage**, is introduced to display the available label space resource in RPD and also the applications that use the label space in RPD. Using this command, the administrator can monitor the available labels in each label space and the applications that are using the labels.

[See [show mpls label usage.](#)]

## Port Security

---

- **Media Access Control Security (MACsec) support (EX9200 switches)**—Starting with Junos OS Release 15.1R1, MACsec is supported on all SFP interfaces on the EX9200-40F-M line card when it is installed in an EX9200 switch. MACsec is an industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is capable of identifying and preventing most security threats, and can be used in combination with other security protocols to provide end-to-end network security. MACsec can only be enabled on domestic versions of Junos OS software. MACsec is standardized in IEEE 802.1AE.

[See [Understanding Media Access Control Security \(MACsec\).](#)]

- **MAC move limiting support (EX9200 switches)**—Starting with Junos OS Release 15.1R1, MAC move limiting is supported on EX9200 switches. MAC move limiting provides port security by controlling the number of MAC address moves that are allowed in a VLAN in one second. When MAC move limiting is configured, the switch tracks MAC address movements on access and trunk interfaces. A MAC address move occurs when an interface on the switch receives a packet with a source MAC address that has already been learned by the switch, but on a different interface. If a MAC address moves more than the configured number of times within one second, you can configure an action to be taken on incoming packets with new source MAC addresses. The incoming packets can be dropped, logged or ignored. You can also specify an action to shutdown or temporarily disable the interfaces associated with that MAC address.

[See [Understanding MAC Limiting and MAC Move Limiting for Port Security on EX Series Switches.](#)]

## Software Installation and Upgrade

---

- **Support for FreeBSD 10 kernel for Junos OS (EX9200 switches)**—Starting with Junos OS Release 15.1, on EX9200, FreeBSD 10 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. Because installation restructures the file system, logs and configurations are lost unless precautions are taken. There are now Junos OS and OAM volumes, which provide the ability to boot from the OAM volume upon failures. Some system commands display a different output than on earlier releases and a few others are deprecated.

[See [Understanding Junos OS with Upgraded FreeBSD.](#)]

## Spanning-Tree Protocols

---

- **Global configuration of spanning-tree protocols (EX Series)**—Starting with Junos OS Release 15.1R1, global configuration of the spanning-tree protocols RSTP, MSTP, and VSTP is supported on EX Series switches with Enhanced Layer 2 Software (ELS) configuration style.

In earlier releases, the ELS software supported configuration of spanning-tree protocols on individual interfaces or on a range of interfaces. It did not support configuration of spanning-tree protocols on all interfaces or disabling spanning-tree protocols on specific interfaces.

Starting with this release, CLI changes in the ELS software provide the options of configuring spanning-tree protocols on all interfaces, disabling the configuration for individual interfaces, and configuring VSTP on all VLANs or on a VLAN group.

[See [Configuring RSTP \(CLI Procedure\)](#), [Configuring MSTP](#), and [Configuring VLAN Spanning-Tree Protocol](#).]

### Related Documentation

- [Changes in Behavior and Syntax on page 10](#)
- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## 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 15.1R2 for the EX Series.

- [Dynamic Host Configuration Protocol on page 11](#)

## Dynamic Host Configuration Protocol

- **Format change for DHCP Option 18**—On EX9200 switches with DHCP snooping configured, when the VLAN ID is appended to the prefix of DHCP option 18, it appears in decimal format instead of hexadecimal format.

### Related Documentation

- [New and Changed Features on page 6](#)
- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R2 for the EX Series.

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

- [Authentication and Access Control](#)
- [J-Web](#)
- [Port Security](#)
- [Spanning-Tree Protocols](#)
- [Virtual Chassis](#)

### Authentication and Access Control

- On EX9200 switches, if you configure a firewall filter such that the number of characters in the filter name, term name, and counter name added together exceeds 128 characters, 802.1X (dot1x) authentication might fail and cause the Network Processing Card (NPC) to crash. As a workaround, configure the filter name, term name, and counter name such that when the sum of the number of characters in those three names is added to the sum of the number of characters in the interface name and the MAC address, the total does not exceed 128. [PR1083132](#)
- On EX9200 switches, 802.1X (dot1x) authentication might not be performed if a voice VLAN is changed or modified to a data VLAN after a client is authenticated in that voice VLAN. This problem occurs when a VoIP VLAN is configured, a client is authenticated in a configured data VLAN, and then the VoIP VLAN is configured as a new data VLAN (that is, you delete the VoIP configuration and delete the current data VLAN membership, and configure the original VoIP VLAN as the new data VLAN). [PR1074668](#)

## J-Web

---

- In the J-Web interface, you cannot commit some of the configuration changes in the Port Configuration page or the VLAN Configuration page because of the following limitations for port-mirroring ports and port-mirroring VLANs:
  - A port configured as the output port for an analyzer cannot be a member of any VLAN other than the default VLAN.
  - A VLAN configured to receive analyzer output can be associated with only one interface.

[PR400814](#)

## Port Security

---

- On EX9200 switches, a DHCPv6 security dynamic entry binding might not work properly on an IPv6 IRB interface that is linked to a DHCP snooping VLAN. [PR1059623](#)
- On EX2200 switches, if you issue the **request system services dhcp release interface-name** operational command, an IP address release message DHCP packet is sent from the client and processed at the server. When the client clears the IP address on the same interface, the kernel generates an event message, which is processed at the client and triggers the DHCP client state machine, which leads to the interface acquiring a new IP address from the server. If you then issue the **show system services dhcp client interface-name** command, the output of that command indicates that the issued **request system services dhcp release interface-name** operational command had no impact.

[PR1072319](#)

## Spanning-Tree Protocols

---

- On an EX9200 switch, an aggregated Ethernet (ae) interface might go down if you configure the **bpdu-block-on-edge** statement in a VSTP configuration. [PR1089217](#)

## Virtual Chassis

---

- On an EX9200 Virtual Chassis, if you restart an FPC with Virtual Chassis ports (VCPs) and there are no other FPCs with VCPs, a Virtual Chassis split might occur and the backup FPC might show a machine check exception and create a Network Processing Card (NPC) core file. [PR1083965](#)

## Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 10](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R2 for the EX Series.

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

- [Dynamic Host Configuration Protocol on page 13](#)
- [Infrastructure on page 13](#)
- [Interfaces and Chassis on page 14](#)
- [J-Web on page 14](#)
- [Software Installation and Upgrade on page 14](#)

### Dynamic Host Configuration Protocol

- On EX9200 switches, when DHCP relay is configured with the DHCP server and DHCP client in separate routing instances, unicast DHCP reply packets, for example, DHCPACK in response to a lease renewal request, might be dropped. [PR1079980](#)
- On EX9200 switches, DHCP snooping and related access security features ARP inspection, IP source guard, Neighbor Discovery inspection and IPv6 source guard, are not supported at the `[edit logical-systems logical-system-name vlans vlan-name forwarding-options dhcp-security]` hierarchy level. [PR1087680](#)

### Infrastructure

- On EX2200 switches, system log messages might display IP addresses in reverse order. For example, an ICMP packet from 10.0.1.114 to 10.0.0.7 might be shown in the log as: `PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 114.1.0.10 7.0.0.10 0 0 (1 packet)`. The correct log message is: `PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 10.0.1.114 10.0.0.7 0 0 (1 packet)`. [PR898175](#)
- On EX9200 switches, the value for the `udpOutDatagrams` object displayed in the output of the `show snmp mib walk decimal udpOutDatagrams` command is different compared with the value for the same object in the output of the `show system statistics udp member 0` command. The value for the `datagrams dropped due to no socket` field is incorrectly used as the `udpOutDatagrams` value in the output for `show snmp mib walk decimal udpOutDatagrams`. As a workaround, use the `show system statistics udp member 0` command. [PR1104831](#)

## Interfaces and Chassis

---

- On EX9200 switches, traffic loss of more than one second (2-6 seconds) might occur on the active node of an MC-LAG when the ICCP (Inter-Chassis Control Protocol) goes down and comes back up. [PR1107001](#)

## J-Web

---

- If you access the J-Web interface using Microsoft Internet Explorer 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, in the Trace Options tab), even though the flags are not configured. As a workaround, use the Mozilla Firefox browser. [PR603669](#)
- On the J-Web interface, on the Route Information page (Monitor > Routing > Route Information), the Next Hop column displays only the interface address, and the corresponding IP address is missing. The title of the first column displays **Static Route Address** instead of **Destination Address**. As a workaround, use the **show route detail CLI** command to fetch the IP address of the next-hop interface. [PR684552](#)
- On the J-Web interface, HTTPS access might work with an invalid certificate. As a workaround, change the certificate and then issue the **restart web-management** command to restart the J-Web interface. [PR700135](#)
- On EX2200-C switches, if you change the media type of an uplink port and commit the change, the Ports Configuration page (Configure > Interfaces > Ports) might not list that uplink port. [PR742847](#)
- If either a copper uplink port or a fiber uplink port is connected on an EX2200-C switch, both might appear to be up in the J-Web dashboard. [PR862411](#)

## Software Installation and Upgrade

---

- On a mixed EX4200 and EX4500 Virtual Chassis or on an EX3300 Virtual Chassis, or EX6200/EX8200, during a Nonstop Software Upgrade (NSSU), packets might be duplicated. [PR1062944](#)
- Substantial traffic losses might occur during an NSSU upgrade on EX4200 and EX4500 Virtual Chassis, EX6200 and EX8200 switches, or EX8200 Virtual Chassis. [PR1062960](#)
- On an EX8200 Virtual Chassis, an NSSU to Release 15.1R1 might fail after the image is pushed to the backup Routing Engine, and a vmcore might be created. [PR1075232](#)
- In Junos Space, the Junos OS Release 15.1R1 image for EX9200 switches is not mapped to the correct platform. As a workaround, in Junos Space, right-click the device image, and select **ex-92xx** in **Modify device image**. [PR1090863](#)
- On EX9200 switches, unified ISSU does not work properly for upgrading to Junos OS Release 15.1R1. Junos Space triggers the upgrades and the upgrades fail. [PR1091610](#)

- On EX9200 switches, during ISSU upgrade from 15.1R1 to 15.1R2, BGP and L3 multicast traffic might be dropped for approximately 30 seconds. [PR1116299](#)
- On EX8200 switches, an NSSU from Junos OS Release 15.1R1 to 15.1R2 fails with the following message: **mgd: unable to execute /var/etc/reboot.ex: Authentication error.** [PR1122628](#)

#### Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 10](#)
- [Known Behavior on page 11](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## 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 online [Junos Problem Report Search](#) application.

- [Resolved Issues: Release 15.1R2 on page 15](#)

#### Resolved Issues: Release 15.1R2

---

- [Dynamic Host Configuration Protocol](#)
- [Class of Service \(CoS\)](#)
- [Interfaces and Chassis](#)
- [Media Access Control Security \(MACsec\)](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Spanning-Tree Protocols](#)
- [VPLS](#)

#### *Dynamic Host Configuration Protocol*

- On EX9200 switches, when DHCP relay is configured using the **forward-only** and **forward-only-replies** statements at the **[edit forwarding-options dhcp-relay]** hierarchy level, if the DHCP local server is also configured with the **forward-snooped-clients** statement at the **[edit system services dhcp-local-server]** hierarchy level, the configuration for **forward-snooped-clients** takes precedence over the configuration for **forward-only** and **forward-only-replies**. As a result, DHCP message exchange between VRFs might not work as expected. [PR1077016](#)

- On EX Series switches except EX9200, the configuration of options for the **circuit-id** CLI statement at the **[edit forwarding-options dhcp-relay group group-name relay-option-82]** hierarchy level does not work as expected. The format of the DHCP option 82 Circuit ID must be **switch-name:physical-interface-name:vlan-name**, but instead, the format is **switch-name:vlan-name**. [PR1081246](#)
- On EX Series switches except EX9200 switches, with DHCP relay configured on the IRB interface for BOOTP relay, if the client is connected to the physical interface that belongs to the same VLAN as the IRB interface, and sends BOOTP request packets to the server, BOOTP reply packets from the server might be dropped on the IRB interface. [PR1096560](#)

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

- On EX4200 switches, if CoS scheduler maps are configured on all interfaces with the **loss-priority** value set to **high**, traffic between different PFEs might be dropped. [PR1071361](#)

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

- On EX9200 switches, if an interface range is configured that includes large-scale physical interfaces, and with the **family** option set to **ethernet-switching**, the configuration might take a long time to commit. [PR1072147](#)
- On EX9200 switches, if an interface for which the MAC move limit action is set to **shutdown** goes down and comes up, and then a Layer 2 learning (l2ald) process restarts, the logical interface remains down even if you issue the command **clear ethernet-switching recovery-timeout**. [PR1072358](#)
- On EX9200 switches, when a MAC move limit is configured on two VLAN members and the limit is configured with the action **vlan-member-shutdown** on two VLAN members, if the limit is reached on one VLAN member, both members are disabled, blocking all traffic. [PR1078676](#)
- On EX9200 platforms, if you configure an MC-LAG with two devices, and then delete and re-create an MC-AE interface, broadcast and multicast traffic that is flooded might loop for several milliseconds. [PR1082775](#)
- An EX9200-40F-M line card drops all traffic on an IRB logical interface, including both data plane and control plane traffic. If an IRB logical interface is configured on an EX9200-40F-M line card as part of a VLAN, any device connected through that interface cannot use Layer 3 forwarding outside the subnet, because the EX9200-40F-M line card does not handle the ARP function correctly. Configuring static ARP on devices using the EX9200 as a gateway is not a workaround, because packets are still dropped if the Routing Engine of the EX9200 has the routes and ARP entry for the destination IP. [PR1086790](#)



### **Media Access Control Security (MACsec)**

- On EX4200 and EX4550 switches, if MACsec is configured to transit traffic between switches through Ethernet over SONET, packets might be dropped. [PR1056790](#)

### **Network Management and Monitoring**

- On EX Series switches, configuring an invalid SNMP source address might prevent SNMP traps from being generated, even after the configuration is corrected with a valid SNMP source address. [PR1099802](#)

### **Platform and Infrastructure**

- On EX4500 and EX4550 switches, if an interface on the EX-SFP-10GE-LR uplink module is disabled by using the CLI command **set interface disable**, and the interface through which a peer device is connected to the interface on the uplink module goes down, CPU utilization of the chassis manager process (chassism) might spike, causing the chassism process to generate a core file. [PR1032818](#)
- On EX Series switches, BFD packets might be sent to a remote neighbor at a rate that exceeds the remote minimum receive interval value. [PR1055830](#)
- On an EX8200 Virtual Chassis, if **vlan-tagging** is configured without specifying the interface family, the PFE might program the local chassis MAC address instead of the router MAC address, which is used for routing. As a workaround, configure family **inet** on the interface. [PR1060148](#)
- On EX Series switches except EX9200 switches, when configuring large numbers of inet addresses on the switch, for example, more than 1000 IP addresses, gratuitous ARP packets might not be sent to peer devices. [PR1062460](#)
- On EX8200 Virtual Chassis, local ECMP hashing changes when a remote (nonlocal) interface flaps if the number of local interfaces does not equal the number of remote interfaces. This might impact ECMP load balancing. [PR1084982](#)
- On EX8200 switches, when the PIM mode is changed between sparse mode and dense mode, the pfem process might generate a core file. [PR1087730](#)
- On EX9200 switches operating in a routing domain with a PIM-embedded IPv6 rendezvous point (RP), accessing the RP after the memory is freed might cause the routing protocol process to generate a core file. [PR1101377](#)

### **Spanning-Tree Protocols**

- On EX Series Virtual Chassis, if STP is configured, and each member's mastership priority values are different, rebooting some or all of the Virtual Chassis members might cause a traffic failure, even after the reboot has completed. [PR1066897](#)
- On EX Series switches except EX9200, when MSTP is configured, the Ethernet switching process (eswd) might generate multiple types of core files in the large-scale VLANs that are associated with multiple spanning-tree Instances (MSTIs). [PR1083395](#)

### VPLS

- On EX9200 switches, when you add a VLAN on an existing virtual-switch instance for virtual private LAN service (VPLS), the label-switched interface (LSI) might not be associated with the new VLAN. [PR1088541](#)

#### Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 10](#)
- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## Documentation Updates

There are no errata or changes in Junos OS Release 15.1R2 for the EX Series switches documentation.

#### Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 10](#)
- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)
- [Product Compatibility on page 19](#)

## Migration, Upgrade, and Downgrade Instructions

This section contains the upgrade and downgrade support policy for Junos OS for the EX Series. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 18](#)

### 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 earlier or later, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

For information about software installation and upgrade, see the [Installation and Upgrade Guide](#).

#### Related Documentation

- [New and Changed Features on page 6](#)
- [Changes in Behavior and Syntax on page 10](#)
- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Product Compatibility on page 19](#)

## Product Compatibility

- [Hardware Compatibility on page 19](#)

### Hardware Compatibility

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

To determine the features supported on EX Series switches 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 6](#)
- [Changes in Behavior and Syntax on page 10](#)

- [Known Behavior on page 11](#)
- [Known Issues on page 13](#)
- [Resolved Issues on page 15](#)
- [Documentation Updates on page 18](#)
- [Migration, Upgrade, and Downgrade Instructions on page 18](#)

## 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 15.1R2 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 21](#)
- [Changes in Behavior and Syntax on page 52](#)
- [Known Behavior on page 67](#)
- [Known Issues on page 68](#)
- [Resolved Issues on page 76](#)
- [Documentation Updates on page 99](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)
- [Product Compatibility on page 113](#)

### New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R2 for the M Series, MX Series, and T Series.

- [Hardware on page 22](#)
- [Bridging and Learning on page 23](#)
- [Class of Service \(CoS\) on page 23](#)
- [High Availability \(HA\) and Resiliency on page 24](#)
- [Interfaces and Chassis on page 26](#)
- [IPv6 on page 30](#)
- [Junos OS XML API and Scripting on page 30](#)
- [Layer 2 Features on page 30](#)
- [Management on page 31](#)
- [MPLS on page 32](#)
- [Multicast on page 33](#)
- [Network Management and Monitoring on page 34](#)
- [Routing Policy and Firewall Filters on page 36](#)

- [Routing Protocols on page 37](#)
- [Services Applications on page 39](#)
- [Software Defined Networking on page 43](#)
- [Software Installation and Upgrade on page 43](#)
- [Subscriber Management and Services \(MX Series\) on page 44](#)
- [User Interface and Configuration on page 50](#)
- [VPNs on page 50](#)

## Hardware

- **New MPC variants that support higher scale and bandwidth (MX Series)**—Starting with Junos OS Release 15.1, the following variants of a new MPC with higher scale and bandwidth are supported on MX Series routers:

- MPC2E-3D-NG—80 Gbps capacity without hierarchical quality of service (HQoS)
- MPC2E-3D-NG-Q—80 Gbps capacity with HQoS
- MPC3E-3D-NG—130 Gbps capacity without HQoS
- MPC3E-3D-NG-Q—130 Gbps capacity with HQoS

The HQoS variants of this MPC support flexible queuing at 80 Gbps or 130 Gbps. See [MIC/MPC Compatibility](#) for supported MICs on these MPCs.



**NOTE:** The MPC2E-3D-NG, MPC2E-3D-NG-Q, MPC3E-3D-NG, and MPC3E-3D-NG-Q are also supported in Junos OS Release 14.1R4. To support these MPCs in 14.1R4, you must install Junos Continuity software. See [Junos Continuity Software](#) for more details.



**NOTE:** The non-HQoS MPCs support MIC-3D-4COC3-1COC12-CE, MIC-3D-8CHOC3-4CHOC12, and MIC-3D-4CHOC3-2CHOC12 when they are upgraded to the HQoS model through a license.

MPC2E-3D-NG and MPC2E-3D-NG-Q do not support MIC3-3D-10XGE-SFPP, MIC3-3D-1X100GE-CFP, MIC3-3D-1X100GE-CXP, and MIC3-3D-2X40GE-QSFPP.

- Starting in Junos OS Release 15.1R1, the Juniper Networks MX2010 and Juniper Networks MX2020 routers support the following new power distribution modules:
  - 7-feed single-phase AC PDM
  - 9-feed single-phase AC PDM
  - 7-feed DC PDM

In addition, this release supports a new optimized power fan tray.

## Bridging and Learning

- **Support for modifying MAC table aging timer for bridge domains (MX Series)**—Starting with Junos OS Release 15.1, you can modify the aging timer for MAC table entries of a bridge domain. When the aging timer for a MAC address in a MAC table expires, the MAC address is removed from the table. This aging process ensures that the router tracks only active MAC addresses on the network and that it is able to flush out MAC addresses that are no longer available.

The default aging timer for MAC entries is 300 seconds. Depending on how long you want to keep a MAC address in a MAC table before it expires, you can either increase or decrease the aging timer. To modify the aging timer for MAC entries in a MAC table, use the **mac-table-aging-timer** statement at one of the following hierarchy levels:

- [edit bridge-domains *bridge-domain-name* bridge-options]
- [edit routing-instances *routing-instance-name* protocols vpls]
- [edit routing-instances *routing-instance-name* protocols evpn]
- **Support for L2TP drain (MX Series)**—Starting in Junos OS Release 15.1, you can prevent the creation of new Layer 2 Tunneling Protocol (L2TP) sessions, destinations, and tunnels at an LNS or a LAC for administrative purposes.

To configure this feature, use the **drain** statement at the [edit services l2tp] hierarchy level. You can configure this feature at the global level or for a specific destination or tunnel. Configuring this feature on a router sets the administrative state of the L2TP session, destination, or tunnel to drain, which ensures that no new destinations, sessions, or tunnels are created at the specified LNS or LAC.



**NOTE:** This feature does not affect existing L2TP sessions, destinations, or tunnels.

[See [Configuring L2TP Drain](#), [show services l2tp destination](#), and [show services l2tp tunnel](#).]

## Class of Service (CoS)

- **Extended MPC support for per-unit schedulers (MX Series)**—Starting in Junos OS Release 15.1 you can configure per-unit schedulers on the non-queuing MPC6E, meaning you can include the **per-unit-scheduler** statement at the [edit interfaces *interface name*] hierarchy level. When per-unit schedulers are enabled, you can define dedicated schedulers for logical interfaces.

Enabling per-unit schedulers on the MPC6E adds additional output to the **show interfaces *interface name* [detail | extensive]** command. This additional output lists the maximum resources available and the number of configured resources for schedulers.

[See [Scheduler Maps and Shaping Rate to DLCIs and VLANs](#).]

- **Change to CoS shaping rate fallback behavior (MX Series)**—Starting in Junos OS Release 15.1, when a CoS service profile is deactivated, the traffic shaping rate falls

back in the following order: ANCP shaping rate, PPPoE IA tag rate, or shaping rate configured in the traffic control profile. In earlier releases, the traffic shaping rate falls back to the ANCP adjusted rate or the traffic control profile value.

Now when an ANCP shaping rate adjustment is removed, the rate falls back to the PPPoE IA tag rate or the traffic control profile value. In earlier releases, the rate falls back to the traffic control profile value.

[See [CoS Adjustment Control Profiles Overview](#).]

- **Hierarchical CoS support for GRE tunnel interface output queues (MX Series routers with MPC5E)**—Starting with Junos OS Release 15.1R2, you can manage output queuing of traffic entering GRE tunnel interfaces hosted on MPC5E line cards in MX Series routers. Support for the **output-traffic-control-profile** configuration statement, which applies an output traffic scheduling and shaping profile to the interface, is extended to GRE tunnel physical and logical interfaces. Support for the **output-traffic-control-profile-remaining** configuration statement, which applies an output traffic scheduling and shaping profile for remaining traffic to the interface, is extended to GRE tunnel physical interfaces.



**NOTE:** Interface sets (sets of interfaces used to configure hierarchical CoS schedulers on supported Ethernet interfaces) are not supported on GRE tunnel interfaces.

[See [Configuring Traffic Control Profiles for Shared Scheduling and Shaping](#).]

---

## High Availability (HA) and Resiliency

- **MX Series Virtual Chassis support for MX2010 and MX2020 member routers (MX Series routers with MPCs/MICs)**—Starting in Junos OS Release 15.1, you can configure an MX2010 router or MX2020 router as a member router in an MX Series Virtual Chassis. In earlier releases, MX2010 routers and MX2020 routers cannot function as member routers in an MX Series Virtual Chassis.

In a two-member Virtual Chassis configuration, the following member router combinations are supported with an MX2010 router or MX2020 router:

- MX960 router and MX2010 router
- MX960 router and MX2020 router
- MX2010 router and MX2020 router
- MX2010 router and MX2010 router
- MX2020 router and MX2020 router

To ensure that a Virtual Chassis configuration consisting of an MX2020 router and *either* an MX960 router or MX2010 router forms properly, you must issue the **request virtual-chassis member-id set member member-id slots-per-chassis slot-count** command, where **member-id** is the member ID (0 or 1) configured for the MX960 router or MX2010 router, and **slot-count** is 20 to match the slot count for the MX2020 router. In addition,



for a Virtual Chassis that includes an MX2020 member router, all four Routing Engines in the Virtual Chassis configuration must have at least 16 gigabytes of memory.

[See [Configuring an MX2020 Member Router in an Existing MX Series Virtual Chassis](#).]

- **Relay daemon code removed for MX Series Virtual Chassis (MX Series routers with MPCs/MICs)**—Starting in Junos OS Release 15.1, the code associated with the relay software process (relayd) has been removed for use with MX Series Virtual Chassis configurations. In earlier releases, the relayd functionality was disabled, but the code implementing this functionality was still present in the software. Removing the relayd functionality and related software code reduces the risk of timing issues for MX Series Virtual Chassis configurations and improves overall performance and stability.

With the removal of the relay daemon code for MX Series Virtual Chassis, certain operational commands no longer display information pertaining to the relayd process in the output for an MX Series Virtual Chassis. Examples of the affected commands include **show system core-dumps**, **show system memory**, and **show system processes**.

In addition, the following relayd error messages have been removed from the software for MX Series Virtual Chassis:

- RELAYD\_COMMAND\_OPTIONS
- RELAYD\_COMMAND\_OPTION\_ERROR
- RELAYD\_SYSCALL\_ERROR
- **Configuration support for multiple MEPs for interfaces belonging to a single VPLS service, CCC, or bridge domain (MX Series)**—Starting with Junos OS Release 15.1, you can configure multiple maintenance endpoints (MEPs) for a single combination of maintenance association and maintenance domain IDs for interfaces belonging to a particular VPLS service, circuit cross-connect (CCC), or bridge domain.  
  
To configure multiple MEPs, use the existing **mep mep-id** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance association ma-name]** hierarchy level.
- **NSR and validation-extension for BGP flowspec**—Starting in Junos OS Release 15.1, changes are implemented to add NSR support for existing inet-flow and inetvpnflow families and to extend routes validation for BGP flowspec. Two new statements are introduced as part of this enhancement.

[See [enforce-first-as](#) and [no-install](#).]

- **Enhancements made to unified ISSU for VRRPv3 to avoid adjacency flap (M Series and MX Series)**—Starting in Junos OS Release 15.1, enhancements have been made to maintain protocol adjacency with peer routers during unified ISSU and to maintain interoperability among equipment and with other Junos OS releases and other Juniper Networks products. This design is for VRRPv3 only. VRRPv1 and VRRPv2 are not supported. The **show vrrp** command output is updated to display unified ISSU information.

[See [show vrrp](#) and [Junos OS Support for VRRPv3](#).]

- **New solution to determine when to tear down old LSP instances (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, a feedback mechanism

supersedes the delay created by using the **optimize-hold-dead-delay** statement. Configure this feature by using the **optimize-adaptive-teardown** statement on routers acting as the ingress for the affected LSPs.

[See [Achieving a Make-Before-Break, Hitless Switchover for LSPs](#), and [optimize-adaptive-teardown](#).]

- **Graceful restart values are configurable at the [edit routing-instances] hierarchy level (M Series and T Series)**—Starting in Junos OS Release 15.1, the **graceful-restart** configuration statement is configurable at the level of individual routing instances. This means you can have different values for different instances. For example, you can have a routing instance configured with IGMP snooping and another with PIM snooping and configure a graceful restart timer value at the instance level that is tuned for each instance.

[See [Configuring Graceful Restart for Multicast Snooping](#) and [graceful-restart \(Multicast Snooping\)](#).]

- **Junos OS achieves higher scaling for VRRP over logical interfaces**—Starting in Junos OS Release 15.1, a new option for the **delegate-processing** statement allows for VRRP over logical interfaces such as aggregated Ethernet and IRB interfaces.

[See [delegate-processing](#).]

---

## Interfaces and Chassis

- **Synchronous Ethernet and Precision Time Protocol (PTP) support on MPC3E-3D-NG-Q, MPC3E-3D-NG, MPC2E-3D-NG-Q, and MPC2E-3D-NG (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1R2, synchronous Ethernet and PTP are supported on MPC3E-3D-NG-Q, MPC3E-3D-NG, MPC2E-3D-NG-Q, and MPC2E-3D-NG. The PTP feature includes support for ordinary clock (OC) and boundary clock (BC).

[See [Precision Time Protocol Overview](#) and [Synchronous Ethernet](#).]

- **VLAN demux support added to MS-DPC (MX Series)**—Starting in Junos OS Release 15.1, the MS-DPC supports VLAN demux interfaces.

[See [Protocols and Applications Supported by the Multiservices DPC \(MS-DPC\)](#).]

- **CFP-100GBASE-ZR (MX Series)**—In Junos OS Release 13.3R6, 14.1R4, 14.2R3, and 15.1R1 and later, the CFP-100GBASE-ZR transceiver provides advanced dual polarization-quadrature phase shift keying (DP-QPSK) coherent digital signal processing (DSP) and forward error correction (FEC)-enabled robust tolerance to optical impairments and supports 80 km reach over single-mode fiber. The transceiver is not specified as part of IEEE 802.3 but is built according to Juniper Networks specifications. The following interface modules support the CFP-100GBASE-ZR transceiver:

- 2x100GE + 8x10GE MPC4E (MPC4E-3D-2CGE-8XGE)
- 100-Gigabit Ethernet MIC with CFP (MIC3-3D-1X100GE-CFP)

For more information about the interface modules, see the “Cables and Connectors” section in the [MX Series Interface Module Reference](#).

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#) and [Supported Network Interface Standards by Transceiver for ACX, M, MX, and T Series Routers](#).]

- **Maximum generation rate for ICMP and ICMPv6 messages is configurable (MX Series)**—Starting in Junos OS Release 15.1R1, you can configure the maximum rate at which ICMP and ICMPv6 messages that are not ttl-expired are generated by using the `icmp` and `icmp6` and configuration statements at the `[edit chassis]` hierarchy level.
- **CPU utilization status (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1, you can view the average CPU utilization status of the local Routing Engine in the past 1 minute, 5 minutes, and 15 minutes using the `show chassis routing-engine` command. You can also view the average CPU utilization status of FPCs in the master Routing Engine in the past 1 minute, 5 minutes, and 15 minutes using the `show chassis fpc` command. In addition, the following three new Juniper Networks enterprise-specific SNMP MIB objects are introduced in the `jnxOperatingTable` table in the `jnxBoxAnatomy` MIB:
  - `jnxOperating1MinAvgCPU`
  - `jnxOperating5MinAvgCPU`
  - `jnxOperating15MinAvgCPU`

[See [jnxBoxAnatomy](#), [show chassis fpc](#), and [show chassis routing engine](#).]

- **Support for a resource-monitoring mechanism using CLI statements and SNMP MIB objects (MX Series routers with DPCs and MPCs)**—Starting in Junos OS Release 15.1, Junos OS supports a resource monitoring capability using both the configuration statements in the CLI and SNMP MIB queries. You can employ this utility to provision sufficient headroom (memory space limits that are set for the application or virtual router) for monitoring the health and operating efficiency of DPCs and MPCs. To configure the resource-monitoring capability on MX240, MX480, MX960, MX2010, and MX2020 routers, include the `resource-monitor` statement and its substatements at the `[edit system services]` hierarchy level. You specify the high threshold value that is common for all the memory spaces or regions and the watermark values for the different memory blocks on DPCs and MPCs.
- **Dynamic learning of source and destination MAC addresses on aggregated Ethernet interfaces (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, support for dynamic learning of the source and destination MAC addresses is extended to aggregated Ethernet interfaces on the following cards: Gigabit Ethernet DPCs on MX Series routers, Gigabit Ethernet IQ and Gigabit Ethernet PICs with SFPs (except the 10-port Gigabit Ethernet PIC and the built-in Gigabit Ethernet port on the M7i router), 100-Gigabit Ethernet Type 5 PIC with CFP configured, and MPC3E, MPC4E, MPC5E, MPC5EQ, and MPC6E MPCs.

[See [Configuring MAC Address Accounting](#).]

- **Support for MACsec (MX Series)**—Starting in Junos OS Release 15.1, you can configure Media Access Control Security (MACsec) on MX Series routers with the enhanced 20-port Gigabit Ethernet MIC (model number MIC-3D-20GE-SFP-E). MACsec is an industry-standard security technology that provides secure communication for almost all types of traffic on Ethernet links. You can enable MACsec using static connectivity association key (CAK) security mode by using the `connectivity-association`

**connectivity-association-name** statement and its substatements at the **[edit security macsec]** hierarchy level. MACsec is supported on MX Series routers with MACsec-capable interfaces. 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.

- **Fabric hardening enhancements (MX Series)**—Starting in Junos OS Release 15.1, fabric hardening can be configured with two new CLI configuration commands, **per fpc bandwidth-degradation** and **per fpc blackhole-action**. Fabric hardening is the process of controlling bandwidth degradation to prevent traffic blackholing. The new commands give you more control over what threshold of bandwidth degradation to react to, and which corrective action to take.

The **per fpc bandwidth-degradation** command determines how the FPC reacts when it reaches a specified bandwidth degradation percentage. The **per fpc bandwidth-degradation** command and the **offline-on-fabric-bandwidth-reduction** commands are mutually exclusive. If both commands are configured, an error is issued during the commit check.

The **per fpc blackhole-action** command determines how the FPC responds to a 100 percent fabric degradation scenario. This command is optional and overrides the default fabric hardening procedures.

- **Support for flexible queuing on non-HQoS MPCs (MX Series)**—Starting in Junos OS Release 15.1, you can enable flexible queuing on non-HQoS MPCs, such as the MPC2E-3D-NG and MPC3E-3D-NG. When flexible queuing is enabled, non-HQoS MPCs support a limited queuing capability of 32,000 queues per slot, including ingress and egress.

You can enable flexible queuing by including the **flexible-queuing-mode** statement at the **[edit chassis fpc]** hierarchy level. When flexible queuing is enabled, the MPC is restarted and is brought online only if the power required for the queuing component is available in the PEM. The MPC remains offline if the PEM cannot meet the power requirement for the queuing component.

The following MICs are supported on non-HQoS MPCs only when flexible queuing is enabled:

- MIC-3D-8CHOC3-4CHOC12
- MIC-3D-4CHOC3-2CHOC12

You must purchase an add-on license to enable flexible queuing on a non-HQoS MPC.

- **Support for dynamic power management (MX Series)**—Starting in Junos OS Release 15.1, MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q support dynamic power management. When you enable dynamic power management, an MPC is powered on only if the power entry module (PEM) can meet the worst-case power requirement for the MPC. Power budgeting for MICs is performed only when a MIC is brought online. Whether or not a new device is powered on depends on the availability of power in the PEM.

You can enable dynamic power management by including the **mic-aware-power-management** statement at the **[edit chassis]** hierarchy level. This

feature is disabled by default. When this feature is disabled, the Chassis Manager checks for the worst-case power requirement of the MICs before allocating power for the MPCs. When dynamic power management is enabled, worst-case power consumption by MICs is not considered while budgeting power for an MPC. Every time you disable or enable dynamic power management, you must restart the chassis or the MPC for the changes to take effect.

- **Maximum generation rate for ICMP and ICMPv6 messages is configurable (MX Series)**—Starting in Junos OS Release 15.1R1, you can configure the maximum rate at which ICMP and ICMPv6 messages that are not ttl-expired are generated by using the `icmp` and `icmp6` and configuration statements at the `[edit chassis]` hierarchy level.
- **Synchronous Ethernet and Precision Time Protocol (PTP) support on MPC4E (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1, synchronous Ethernet and PTP are supported on MPC4E. The PTP feature includes support for ordinary clock (OC) and boundary clock (BC).

[See [Precision Time Protocol Overview](#), [Synchronous Ethernet](#), and [Protocols and Applications Supported by the MX240, MX480, MX960, MX2010, and MX2020 MPC4Es](#).]

- **Support for hyper mode to increase packet processing rate on enhanced MPCs (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 15.1, MPC3E, MPC4E, MPC5E, and MPC6E support the hyper mode feature. Enabling the hyper mode feature increases the rate at which a data packet is processed, which results in the optimization of the lifetime of a data packet. Optimization of the data packet lifetime enables better performance and throughput.



**NOTE:** You can enable hyper mode only if the network-service mode on the router is configured as either `enhanced-ip` or `enhanced-ethernet`. Also, you cannot enable the hyper mode feature for a specific Packet Forwarding Engine on an MPC—that is, when you enable the feature, it is applicable for all Packet Forwarding Engines on the router.

When you enable the hyper mode feature, the following features are not supported:

- Creation of Virtual Chassis.
- Interoperability with legacy DPCs, including MS-DPCs. The MPC in hyper mode accepts and transmits data packets only from other existing MPCs.
- Interoperability with non-Ethernet MICs and non-Ethernet Interfaces such as channelized interfaces, multilink interfaces, and SONET interfaces.
- Padding of Ethernet frames with VLAN.
- Sending Internet Control Message Protocol (ICMP) redirect messages.
- Termination or tunneling of all subscriber-based services.

To configure the hyper mode feature, use the `hyper-mode` statement at the `[edit forwarding-options]` hierarchy level. To view the changed configuration, use the `show forwarding-options hyper-mode` command.

## IPv6

---

- **Support for outbound-SSH connections with IPv6 addresses (M Series, MX Series, and T Series)**—Starting with Release 15.1, Junos OS supports outbound-SSH connections with devices having IPv6 addresses.

[See [outbound-ssh](#), [Configuring Outbound SSH Service](#), and [Establishing an SSH Connection for a NETCONF Session](#).]

## Junos OS XML API and Scripting

---

- **Support for replacing patterns in configuration data within NETCONF and Junos XML protocol sessions (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, you can replace variables and identifiers in the candidate configuration when performing a `<load-configuration>` operation in a Junos XML protocol or NETCONF session. The **replace-pattern** attribute specifies the pattern to replace, the **with** attribute specifies the replacement pattern, and the optional **upto** attribute indicates the number of occurrences to replace. The scope of the replacement is determined by the placement of the attributes in the configuration data. The functionality of the attribute is identical to that of the **replace pattern** configuration mode command in the Junos OS CLI.

[See [Replacing Patterns in Configuration Data Using the NETCONF or Junos XML Protocol](#).]

- **Junos OS SNMP scripts to support custom MIBs (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, you can use Junos OS SNMP scripts to support custom MIBs until they are implemented in Junos OS. SNMP scripts are triggered automatically when the SNMP manager requests information from the SNMP agent for an object identifier (OID) that is mapped to an SNMP script for an unsupported OID. The script acts like an SNMP subagent, and the system sends the return value from the script to the network management system (NMS).

[See [SNMP Scripts Overview](#).]

## Layer 2 Features

---

- **Configuration support for backup liveness detection between multichassis link aggregation peers (MX Series)**—Starting with Junos OS Release 15.1, you can configure backup liveness detection between multichassis link aggregation (MC-LAG) peers.

Backup liveness detection determines the peer status (that is, whether the peer is up or down) by exchanging keepalive messages between two MC-LAG peers on a configured IP address. MC-LAG peers use an Inter-Chassis Control Protocol (ICCP) connection to communicate. When an ICCP connection is operationally down, a peer can send liveness detection requests to determine the peer status. If a peer fails to respond to the liveness detection request within a specified time interval, the liveness detection check fails and the peer is concluded to be down.

To configure backup liveness detection between MC-LAG peers, use the **backup-liveness-detection backup-peer-ip *backup-peer-ip-address*** statement at the **[edit protocols iccp peer]** hierarchy level.

[See [Configuring Multichassis Link Aggregation on MX Series Routers](#) and [show iccp](#).]

- **Support for PTP over Ethernet (MX Series)**—Starting in Junos OS Release 15.1, Precision Time Protocol (PTP) is supported over Ethernet links on MX Series routers. This functionality is supported in compliance with the IEEE 1588-2008 specification.

Some base station vendors might use only packet interfaces using Ethernet encapsulation for PTP for time and phase synchronization. To provide packet-based timing capability to packet interfaces used by such vendors, you can configure Ethernet encapsulation for PTP on the master port of any node (that is, an MX Series router) that is directly connected to the base station.

To configure Ethernet as the encapsulation type for the transport of PTP packets on master or slave interfaces, use the **transport 802.3** statement at the **[edit protocols ptp slave interface *interface-name* multicast-mode]** or **[edit protocols ptp master interface *interface-name* multicast-mode]** hierarchy level.

[See [Configuring Precision Time Protocol](#).]

- **Support extended for Layer 2 features (MX Series routers with MPC5E and MPC6)**—Starting with Junos OS Release 15.2, Junos OS extends support for the following Layer 2 features on MX Series routers with MPC5E and MPC6:
  - Active-active multihoming support for EVPNs
  - Ethernet frame padding with VLAN for DPCs and MPCs
  - IEEE 802.1ad provider bridges
  - IGMP snooping with bridging, IRB, and VPLS
  - Layer 2 and Layer 2.5 integrated routing and bridging (IRB) and Spanning Tree Protocols (xSTP)
  - Layer 2 protocol tunneling (L2PT) support
  - Layer 2 support for MX Series Virtual Chassis
  - Layer 2 Tunneling Protocol (L2TP)
  - Link aggregation group (LAG)—VLAN-CCC encapsulation
  - Loop Detection using the MAC address Move
  - Multichassis LAG—active/active and active/standby
  - Multichassis LAG—active/active with IGMP snooping
  - Truck ports

[See [Layer 2 Overview, Routing Instances, and Basic Services Feature Guide for Routing Devices](#).]

## Management

- **Support for enforcing RFC-compliant behavior in NETCONF sessions (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, you can require that the NETCONF server enforce certain behaviors during the NETCONF session by configuring the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level. When you configure the **rfc-compliant** statement, the NETCONF server explicitly declares the

NETCONF namespace in its replies and qualifies all NETCONF tags with the **nc** prefix. Also, **<get>** and **<get-config>** operations that return no configuration data do not include an empty **<configuration>** element in RPC replies.

[See [Configuring RFC-Compliant NETCONF Sessions](#).]

- **New YANG features including configuration hierarchy must constraints published in YANG and a new module that defines Junos OS YANG extensions (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, the Juniper Networks **configuration** YANG module includes configuration constraints published using either the YANG **must** statement or the Junos OS YANG extension **junos:must**. Constraints that cannot be mapped directly to YANG's **must** statement, which include expressions containing special keywords or symbols such as **all**, **any**, **unique**, **\$**, **\_**, and wildcard characters, are published using **junos:must**.

The new **junos-extension** module contains definitions for Junos OS YANG extensions including the **must** and **must-message** keywords. The **junos-extension** module is bound to the namespace URI <http://yang.juniper.net/yang/1.1/je> and uses the prefix **junos**. You can download Juniper Networks YANG modules from the website, or you can generate the modules by using the **show system schema** operational mode command on the local device.

[See [Using Juniper Networks YANG Modules](#).]

---

## MPLS

- **New command to display the MPLS label availability in RPD (MX Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, a new show command, **show mpls label usage**, is introduced to display the available label space resource in RPD and also the applications that use the label space in RPD. Using this command, the administrator can monitor the available labels in each label space and the applications that are using the labels.

[See [show mpls label usage](#).]

- **Pro-active loss and delay measurement (MX Series routers with MPCs and MICs only)**—Starting in Junos OS Release 15.1, this feature enables you to monitor and measure packet loss, packet delay, or both for associated bidirectional MPLS ultimate-hop popping point-to-point label-switched paths (LSPs), using the **show performance-monitoring mpls lsp** command. This command provides a summary of the performance metrics for packet loss, two-way channel delay and round trip delay, as well as related metric like delay variation and channel throughput.

You can configure pro-active loss and delay measurement using the **performance-monitoring** configuration statement. This functionality provides real-time visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

[See [Configuring Pro-Active Loss and Delay Measurements](#).]

- **Configuring Layer 3 VPN egress protection with PLR as protector (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, this feature addresses a special scenario of egress node protection, where the point of local repair (PLR) and the



protector are co-located as one router. In this case, there is no need to have a bypass LSP reroute traffic during local repair.

In the co-located protector model, the PLR or the protector is directly connected to the CE device through a backup AC, while in the Centralized protector model, the PLR or the protector has an MPLS tunnel to the backup PE device.

[See [Example: Configuring Layer 3 VPN Egress Protection with PLR as Protector.](#)]

- **Support for NSR, IGP-FA, and static route on container LSPs (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, container label-switched paths (LSPs) provide support for nonstop active routing (NSR), IGP forwarding adjacency, and static routes to address the requirements of a wider business case.

NSR synchronizes the LSP state between redundant Routing Engines, thereby reducing the time to rebuild the container LSP upon a Routing Engine switchover and avoiding traffic loss. Because IGP forwarding adjacency and static routes are widely deployed for RSVP point-to-point LSPs, and container LSPs are dynamically created point-to-point LSPs, these features are also required to fully deploy container LSPs in the field.

[See [Example: Configuring Dynamic Bandwidth Management Using Container LSPs.](#)]

## Multicast

- **Latency fairness optimized multicast (MX Series)**—Starting with Junos OS Release 15.1R1, you can reduce latency in the multicast packet delivery by optimizing multicast packets sent to the Packet Forwarding Engines. You can achieve this by enabling the **ingress** or **local-latency-fairness** option in the **multicast-replication** configuration statement at the **[edit forwarding-options]** hierarchy level. The **multicast-replication** statement is supported only on platforms with the **enhanced-ip** mode enabled. This feature is not supported in VPLS networks, point-to-multipoint connections, and on integrated routing and bridging (IRB) interfaces.

[See [multicast-replication.](#)]

- **IGMP snooping on pseudowires (MX Series)**—Starting in Junos OS Release 15.1, you can prevent multicast traffic from traversing a pseudowire (to egress PE routers) unless there are IGMP receivers for the traffic.

The default IGMP snooping implementation for a VPLS instance adds each pseudowire interface to its **oif** list. This includes traffic sent from the ingress PE router to the egress PE router regardless of interest. The **snoop-pseudowires** option prevents multicast traffic from traversing the pseudowire (to the egress PE routers) unless there are IGMP receivers for the traffic. In other words, multicast traffic is forwarded only to VPLS core interfaces that are either router interfaces or IGMP receivers. In addition to the benefit of sending traffic to interested PE routers only, **snoop-pseudowires** optimizes a common path between PE-P routers wherever possible. Thus, if two PE routers connect through the same P router, only one copy of the packet is sent because the packet is replicated on only those P routers for which the path is divergent.

[See [snoop-pseudowires.](#)]

- **Sender-based RPF and hot-root standby for ingress replication provider tunnels (MX Series routers with MPCs running in "enhanced-ip" mode)**—Starting in Junos OS Release 15.1, support has been added for sender-based RPF and hot-root standby to ingress replication for selective (not inclusive) provider tunnels. This feature extends the sender-based RPF functionality for RSVP-P2MP added in Junos OS Release 14.2, which, in conjunction with hot-root standby, provides support for live-live NGEN MVPN traffic. The configuration of the router, whether for RSVP-P2MP or ingress replication provider tunnels, determines the form of sender-based RPF and hot-root standby that are implemented when their respective CLI configurations are enabled.

Ingress replication works by introducing a unique VPN label to advertise each upstream PE router per VRF. This allows the ingress replication to distinguish the sending PE router and the VRF. When ingress replication is used as the selective provider tunnel, ingress replication tunnels must also be configured for all interested egress PE routers or border routers. When sender-based RPF is disabled, it causes all type 4 routes to be re-advertised with the VT/LSI label. Ingress replication is not intended to work in S-PMSI only configurations.

[See [hot-root-standby \(MBGP MVPN\)](#) and [sender-based-rpf \(MBGP MVPN\)](#).]

- **Fast-failover according to flow rate (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, for routers operating in Enhanced IP Network Services mode, you can configure a threshold that triggers fast failover in NG MVPNs with hot-root standby on the basis of aggregate flow rate. For example, fast failover (as defined in *Draft Morin L3VPN Fast Failover 05*) is triggered if the flow rate of monitored multicast traffic from the provider tunnel drops below the set threshold.

[See [sender-based-rpf \(MBGP MVPN\)](#).]

---

## Network Management and Monitoring

- **Configuring SNMP to match jnxNatObjects values for MS-DPC and MS-MIC (MX Series)**—In Junos OS Release 13.3R7, 14.1R6, 14.2R4, and 15.1R2, you can configure the `snmp-value-match-msmic` statement at the `[edit services service-set service-set-name nat-options]` hierarchy level.

In networks where both MS-DPC and MS-MIC are deployed, you can configure this statement to ensure that the values for MS-MIC-specific objects in the `jnxNatObjects` MIB table match the values for MS-DPC objects. By default, this feature is disabled. You can use the `deactivate services service-set service-set-name nat-options snmp-value-match-msmic` configuration mode command to disable this feature.

- **Tracing tacplus processing (M Series, MX Series, and T Series)**—Starting in Release 15.1, Junos OS allows users to trace tacplus processing. To trace tacplus processing, include the `tacplus` statement at the `[edit system accounting traceoptions flag]` hierarchy level.

[See [traceoptions \(System Accounting\)](#).]

- **Support for multi-lane digital optical monitoring (DOM) MIB (MX960, MX480, and MX240)**—Starting with Release 15.1, Junos OS supports the following SNMP tables and objects in the `jnxDomMib` MIB that gives you information about multi-lane digital optical modules in 10-gigabit small form-factor pluggable transceiver (XFP), small

formfactor pluggable transceiver (SFP), small form-factor pluggable plus transceiver (SFP+), quad small form-factor pluggable transceiver (QSFP), and C form-factor pluggable transceiver (CFP):

- `jnxDomModuleLaneTable`
- `jnxDomCurrentModuleVoltage` in `jnxDomCurrentTable` table
- `jnxDomCurrentModuleVoltageHighAlarmThreshold` in `jnxDomCurrentTable` table
- `jnxDomCurrentModuleVoltageLowAlarmThreshold` in `jnxDomCurrentTable` table
- `jnxDomCurrentModuleVoltageHighWarningThreshold` in `jnxDomCurrentTable` table
- `jnxDomCurrentModuleVoltageLowWarningThreshold` in `jnxDomCurrentTable` table
- `jnxDomCurrentModuleLaneCount` in `jnxDomCurrentTable`

Junos OS also supports the `jnxDomLaneNotifications` traps.

[See [Enterprise-Specific SNMP Traps Supported by Junos OS](#), and [Digital Optical Monitoring MIB](#).]

- **SNMP support for Service OAM (SOAM) performance monitoring functions (MX Series)**—Starting in Junos OS Release 15.1, SNMP supports Service OAM (SOAM) performance monitoring functions that are defined in Technical Specification MEF 17, the Service OAM performance monitoring requirements specified in SOAM-PM, and the Service OAM management objects specified in Technical Specification MEF 7.1.

A new enterprise-specific MIB, SOAM PM MIB, that defines the management objects for Ethernet services operations, administration, and maintenance for performance monitoring, has been added and SNMP support is available for the MIB objects defined in Technical Specification MEF 36.

- **SNMP support for fabric and WAN queue depth monitoring (MX Series)**—Starting in Release 15.1, Junos OS supports monitoring of fabric and WAN queues at the Packet Forwarding Engine level. You can configure fabric and WAN queue depth monitoring by enabling the `queue-threshold` statement at the `[edit chassis fpc slot-number traffic-manager]` hierarchy level. When the `fabric-queue` and `wan-queue` statements are configured, an SNMP trap is generated when the fabric queue or WAN queue depth exceeds the configured threshold value.

The SNMP traps `jnxCosFabricQueueOverflow`, `jnxCosFabricQueueOverflowCleared`, `jnxCosWanQueueOverflow`, and `jnxCosWanQueueOverflowCleared` have been added to the Juniper Networks enterprise-specific Class of Service (COS) MIB to support fabric and WAN queue monitoring.

- **SNMP support for monitoring fabric power utilization (MX Series)**—Starting in Release 15.1, Junos OS supports monitoring of fabric power utilization. An SNMP trap is generated whenever the fabric power consumption exceeds the configured threshold value. The SNMP trap `jnxFabricHighPower` has been added to the `jnxFabricChassisTraps` group to indicate excessive power consumption. The SNMP trap `jnxFabricHighPowerCleared`

added to the `jnxFabricChassisOKTraps` group sends notification when the condition of consuming excessive power is cleared.

- **Support for the interface-set SNMP index (MX Series)**—Starting with Release 15.1R2, Junos OS supports the interface-set SNMP index that provides information about interface-set queue statistics. The following interface-set SNMP index MIBs are introduced in the Juniper Networks enterprise-specific Class-of-Service MIB:
  - `jnxCosIfTable` in `jnxCos` MIB
  - `jnxCosIfsetQstatTable` in `jnxCos` MIB

[See [jnxCosIfTable](#) and [jnxCosIfsetQstatTable](#).]

---

## Routing Policy and Firewall Filters

- **Support for consistent load balancing for ECMP groups (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, on MX Series routers with modular port concentrators (MPCs) only, you can prevent the reordering of flows to active paths in an ECMP group when one or more paths fail. Only flows that are inactive are redirected. This feature applies only to Layer 3 adjacencies learned through external BGP connections. It overrides the default behavior of disrupting all existing, including active, TCP connections when an active path fails. Include the **consistent-hash** statement at the `[edit policy-options policy-statement policy-statement-name then load-balance]` hierarchy level. You must also configure a global per-packet load-balancing policy.

[See [Actions in Routing Policy Terms](#).]

- **New fast-lookup-filter statement (MX240, MX480, MX960, MX2010, and MX2020 routers with MPC5E, MPC5EQ, and MPC6E MPCs and compatible MICs)**—Starting in Junos OS Release 15.1, the **fast-lookup-filter** option is available at the `[edit firewall family (inet | inet6) filter filter-name]` hierarchy level. This allows for hardware assist from compatible MPCs in the firewall filter lookup. There are 4096 hardware filters available for this purpose, each of which can support up to 255 terms. Within the firewall, filters and their terms, ranges, prefix lists, and the except keyword are all supported. Only the inet and inet6 protocol families are supported.

[See [fast-lookup-filter](#).]

- **New forwarding-class-accounting statement (MX Series)**—Starting in Junos OS Release 15.1, you can enable new forwarding class accounting statistics at the `[edit interfaces interface-name]` and `[edit interfaces interface-name unit interface-unit-number]` hierarchy levels. These statistics replace the need to use firewall filters for gathering accounting statistics. Statistics can be gathered in ingress, egress, or both directions. Statistics are displayed for IPv4, IPv6, MPLS, Layer 2, and other families.

[See [forwarding-class-accounting](#).]

## Routing Protocols

- **BGP Prefix Independent Convergence for inet (MX Series routers with MPCs)**—Beginning with Junos OS Release 15.1, BGP Prefix Independent Convergence (PIC), which was initially supported for Layer 3 VPN routers, is extended to BGP with multiple routes in the global tables such as inet and inet6 unicast, and inet and inet6 labeled unicast. When the BGP PIC feature is enabled on a router, BGP installs to the Packet Forwarding Engine the second best path in addition to the calculated best path to a destination. When an IGP loses reachability to a prefix, the router uses this backup path to minimize traffic loss until the global convergence through the BGP is resolved, thereby drastically reducing the outage duration.

[See [Use Case for BGP PIC for Inet.](#)]

- **Entropy label support for BGP-LU (MX Series routers with MPCs, and T Series routers with HC-FPC)**—Beginning with Junos OS Release 15.1, entropy labels for BGP labeled unicast LSPs are supported. You can configure entropy labels for BGP labeled unicasts to achieve end-to-end load balancing. BGP labeled unicasts generally concatenate RSVP or LDP LSPs across multiple IGP areas or multiple autonomous systems. RSVP or LDP entropy labels are popped at the penultimate hop node, together with the RSVP or LDP label. However, there are no entropy labels at the stitching points. Therefore, in the absence of entropy labels, the load-balancing decision at the stitching points was based on deep packet inspection. Junos OS now allows the insertion of entropy labels at the BGP labeled unicast LSP ingress.

[See [Entropy Label for BGP Labeled Unicast LSP Overview.](#)]

- **Multi-instance support for RSVP-TE (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 15.1, multi-instance support is extended to the existing MPLS RSVP-Traffic Engineering (TE) functionality. This support is available only for a virtual router instance type. A router can create and participate in multiple independent TE topology partitions, which allows each partitioned TE domain to scale independently.

Multi-instance support is also extended for LDP over RSVP tunneling for a virtual router routing instance. This supports splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

[See [Tunneling LDP LSPs in RSVP LSPs Overview.](#)]

- **Support for long-lived BGP graceful restart (M Series, MX Series, and T Series)**—Starting in Release 15.1, Junos OS supports the mechanism to preserve BGP routing details from a failed BGP peer for a longer period than the duration for which such routing information is maintained using the BGP graceful restart functionality. To enable the BGP long-lived graceful restart capability, include the **long-lived receiver enable** statement at the `[edit protocols bgp graceful-restart]`, `[edit protocols bgp group group-name graceful-restart]`, and `[edit protocols bgp group group-name neighbor neighbor-address graceful-restart]` hierarchy levels.
- **Selection of backup LFA for OSPF routing protocol (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, the default loop-free alternative (LFA) selection algorithm or criteria can be overridden with an LFA policy. These policies are

configured per destination per primary next-hop interface or per destination. These backup policies enforce LFA selection based on admin-group, srlg, node, bandwidth, protection-type, and metric attributes of the backup path. During backup shortest-path-first (SPF) computation, each attribute (both node and link) of the backup path, stored per backup next hop, is accumulated by IGP. For the routes created internally by IGP, the attribute set of every backup path is evaluated against the policy configured per destination per prefix primary next-hop interface. The first or the best backup path is selected and installed as the backup next hop in the routing table.

[See [Example-configuring-backup-selection-policy-for-ospf-protocol.](#)]

- **Remote LFA support for LDP in OSPF (MX Series)**—Beginning with Junos OS Release 15.1, you can configure a remote loop-free alternate (LFA) to extend the backup provided by the LFA in an OSPF network. This feature is useful especially for Layer 1 metro-rings where the remote LFA is not directly connected to the PLR. The existing LDP implemented for the MPLS tunnel setup can be reused for the protection of OSPF networks and subsequent LDP destinations thereby eliminating the need for RSVP-TE backup tunnels for backup coverage.

[See [Example: Configuring Remote LFA Over LDP Tunnels in OSPF Networks.](#)]

- **Configuring per-interface NDP cache protection (MX Series)**—Starting in Junos OS Release 15.1, you can configure the per-interface neighbor discovery process (NDP).

NDP is that part of the control plane that implements Neighbor Discovery Protocol. NDP is responsible for performing address resolution and maintaining the neighbor cache. NDP picks up requests from the shared queue and performs any necessary discovery action.

NDP queue limits can be enforced by restricting queue size and rate of resolution, and prioritizing certain categories of NDP traffic. The queue limits can be enforced through dynamically configurable queue sizes, for which you can tune global and per interface (IFL) limits for configuring system-wide limits on the NDP queue.

[See [Example: Configuring NDP Cache Protection to Prevent Denial-of-Service Attacks.](#)]

- **Configuring per-prefix LFA and node to link protection fallback for OSPF (M Series, MX Series, and T Series)** —Starting in Junos OS Release 15.1, you can configure the following features for OSPF:

- Per-prefix loop-free alternates (LFAs)
- Fallback to link protecting LFA from node protecting LFA

In certain topologies and usage scenarios, it might be possible that multiple destinations are originating the same prefix and there is no viable LFA to the best prefix originator, while a non-best prefix originator has one.

In certain topologies it might be desirable to have local repair protection to node failures in the primary next hop, which might not be available. In that case, to ensure that some level of local repair capabilities exist, a fallback mechanism is required. Since the link protection is less stringent than node protection, it might be possible that link protection exists and provides the same to those destinations (and hence the prefixes originated by the destinations).

[See [Configuring Per-Prefix LFA for OSPF](#) and [Configuring Node to Link Protection Fallback for OSPF](#).]

- **OSPFv3-TTL propagation policy for TE-Shortcuts and FA-LSPs in-line with other modules in the system (MX Series)**—Starting in Junos OS Release 15.1R2, the OSPFv3-TTL propagation policy will be dictated by MPLS-TTL propagation policy which, by default, allows propagation of TTL.

This change makes behavior of OSPFV3 in-line with the default behavior of rest of the system, allowing you to *disable* TTL propagation for the above mentioned LSPs and for traffic-engineering-shortcuts (TE-Shortcuts) and forwarding adjacency LSPs (FA-LSPs) using OSPFv3 as IGP, by configuring the **no-propagate-ttl** statement at the **[edit protocols mpls]** hierarchy.

- **OSPF domain-id interoperability (MX Series)**— Starting in Junos OS Release 15.1R2, to enable interoperability with routers from other vendors, you can set the AS number for **domain-id** attributes to 0 at the following hierarchical levels:

[edit routing-instances *routing-instance name* protocols ospf domain-id]

or

[edit policy-options community *community name* members]



**CAUTION:** Do not downgrade Junos OS after configuring the AS number for domain-id attributes to 0. Set the AS number to a nonzero value and commit the configuration before downgrading Junos OS.

## Services Applications

- **Support for NAT port block allocation (MX Series routers with MS-MPCs or MS-MICs)**—Starting in Junos OS Release 15.1, you can configure port block allocation for NAT with port translation (NAPT) on MX Series routers with MS-MPCs or MS-MICs. The existing CLI and configuration procedures used for other interface cards remain unchanged. Deterministic port block allocation is not supported.

[See [secured-port-block-allocation](#) and [Configuring Address Pools for Network Address Port Translation \(NAPT\) Overview](#).]

- **Support for inline 6rd and 6to4 (MX Series routers with MPCs )**—Starting in Junos OS Release 15.1, you can configure inline 6rd or 6to4 on an MPC. You can use the inline capability to avoid the cost of using MS-DPCs for required tunneling, encapsulation, and decapsulation processes. Anycast is supported for 6to4 using next-hop service interfaces. Hairpinning is also supported for traffic between 6rd domains. The CLI configuration statements for inline and service PIC-based 6rd remain unchanged. To implement the inline functionality, configure service interfaces on the MPC as inline services interfaces (si-) rather than as multiservices (ms-) interfaces. Two new operational mode commands have been added: **show services inline software statistics** and **clear services inline software statistics**.

[See [Tunneling Services for IPv4-to-IPv6 Transition Overview](#), [show services inline software statistics](#), and [clear services inline software statistics](#).]



- **Support for interim logging for NAT port block allocation (MX Series routers with MS-MPCs or MS-MICs)**—Starting in Junos OS Release 15.1, you can configure interim logging for NAT with port translation on MX Series routers with MS-MPCs or MS-MICs. Default logging sends a single log entry for ports allocated to a subscriber. These syslog entries can be lost for long running flows. Interim logging triggers re-sending of logs at configured time intervals for active blocks that have traffic on at least one of the ports of the block, ensuring that there is a recent syslog entry for active blocks. You can specify interim logging by including the **pba-interim-logging-interval** statement at the **[edit interfaces *interface-name* services-options]** hierarchy level.

[See [pba-interim-logging-interval](#) and [Configuring NAT Session Logs](#).]

- **Support for NAT mapping controls and EIF session limits (MX Series routers with MS-MICs)**—Starting in Junos OS Release 15.1, you can control network address translation (NAT) mapping refresh behavior and establish endpoint-independent filtering session limits for flows on MS-MICs. The following features, previously introduced on MS-DPCs, are available:
  - Clear NAT mappings using the **clear services nat mappings** command.
  - Configure criteria for refreshing NAT mappings for inbound flows and outbound flows. To configure refresh criteria, include the **mapping-refresh** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.
  - Configure a limit for inbound sessions for an EIF mapping. To configure this limit, include the **eif-flow-limit** statement at the **[edit services nat rule *rule-name* term *term-name* then translated secure-nat-mapping]** hierarchy level.
  - Configure a limit for the number of dropped flows (ingress, egress, or both) for a specified service set. To configure this limit, include the **max-drop-flows** statement at the **[edit services service-set *service-set-name*]** hierarchy level.

[See [clear-services-nat-mappings](#), [clear-services-nat-flows mapping-refresh](#), [eif-flow-limit](#), and [max-drop-flows](#).]

- **Support for per-service throughput for NAT and inline flow monitoring services (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1, you can configure the capability to transmit the throughput details per service for Junos Address Aware (carrier-grade NAT) and Junos Traffic Vision (previously known as J-Flow) in the last time interval to an external log collector. The default time interval at which the throughput data is sent is 300 seconds, which you can configure to suit your network needs. Multiple instances of the same service running on different PICs within a router are supported. This functionality is supported on MX Series routers with MS-MPCs and MS-MICs, and also in the MX Series Virtual Chassis configuration.
- **Support for generation of SNMP traps and alarms for inline video monitoring (MX Series)**—Starting in Junos OS Release 15.1, SNMP support is introduced for the media delivery index (MDI) metrics of inline video monitoring. Inline video monitoring is available on MX Series routers using only MPCE1, MPCE2, and MPC-16XGE. Until Junos OS Release 14.2, inline MDI generated only syslogs when the computed MDI metric value was not within the configured range. SNMP support is now added to enable SNMP traps to be triggered when the computed delay factor, media rate variation



(MRV), or media loss rate (MLR) values are not within the configured range. You can retrieve the MDI statistics, flow levels, error details, and MDI record-level information using SNMP Get and Get Next requests. The SNMP traps and alarms that are generated when the MDI metrics exceed the configured ranges can be cleared as necessary. Also, you can control the flooding of SNMP traps on the system.

- **Support for Layer 2 services over GRE tunnel interfaces (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, you can configure Layer 2 Ethernet services over GRE interfaces (*gr-fpc/pic/port* to use GRE encapsulation). To enable Layer 2 Ethernet packets to be terminated on GRE tunnels, you must configure the bridge domain protocol family on the gr- interfaces and associate the gr- interfaces with the bridge domain. You must configure the GRE interfaces as core-facing interfaces, and they must be access or trunk interfaces. To configure the bridge domain family on gr- interfaces, include the **family bridge** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level.
- **Support for stateless source IPv6 prefix translation (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, you can configure stateless translation of source address prefixes in IPv6 networks. This capability is supported on MX Series routers with MPCs where inline NAT is supported. To configure stateless network prefix translation for IPv6 packets (NPTv6), include the **translation-type nptv6** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level. The NPTv6 translator translates the source address prefix in such a way that the transport layer checksum of the packet does not need to be recomputed.
- **Support for logging flow monitoring records with version 9 and IPFIX templates for NAT events (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1, you can configure MX Series routers with MS-MPCs and MS-MICs to log NAT events by using Junos Traffic Vision (previously known as J-Flow) version 9 or IPFIX (version 10) template format. NAT event logger generates messages in flow monitoring format for various NAT events, such as the creation of a NAT entry, deletion of a NAT entry, and for invalid NAT processing. These events also support NAT64 translations (translation of IPv6 addresses to IPv4 addresses), binding information base (BIB) events, and more detailed error generation. The generated records or logs for NAT events in flow template format are sent to the specified host or external device that functions as the NetFlow collector.
- **Support for unified ISSU on inline LSQ interfaces (MX Series)**—Starting in Junos OS Release 15.1, unified in-service software upgrade (ISSU) is supported on inline link services intelligent queuing (IQ) (lsq-) interfaces on MX Series routers. Unified ISSU enables an upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. The inline LSQ logical interface (*lsq-slot/pic/0*) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC.
- **Inline TWAMP requester support (MX Series)**—Starting in Junos OS Release 15.1, MX Series routers support the inline Two-Way Active Measurement Protocol (TWAMP) control-client and session-sender for transmission of TWAMP probes using IPv4 between the sender (control-client) and the receiver (session-sender or server). The control-client and session-sender reside on the same router. The TWAMP control-client can also work with a third-party server implementation.

- **Ethernet over generic routing encapsulation (GRE) and GRE key support for label blocks (MX Series)**—Starting in Junos OS Release 15.1, MX Series routers support the following in compliance with RFC 2890:
  - Adding a bridge family on general tunneling protocol
  - Switching functionality supporting connections to the traditional Layer 2 network and VPLS network
  - Routing functionality supporting integrated routing and bridging (IRB)
  - Configuring the GRE key and performing the **hash load balance** operation both at the **gre tunnel initiated** and **transit routers** hierarchies
  - Providing statistics for the GRE-L2 tunnel
- **Support for IRB in a P-VLAN bridge domain (MX Series)**—Starting in Junos OS Release 15.1, MX Series routers support IRB in a private VLAN (P-VLAN) bridge domain. All IP features such as IP multicast, IPv4, IPv6, and VRRP that work for IRB in a normal bridge domain also work for IRB in a P-VLAN bridge domain.
- **Enhancements to the RFC 2544-based benchmarking tests (MX104)**—Starting in Junos OS Release 15.1, MX104 routers support RFC 2544-based benchmarking tests for Ethernet transparent LAN (E-LAN) services configured using LDP-based VPLS and BGP-based VPLS. The RFC 2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before the E-LAN service is activated. The tests measure throughput, latency, frame-loss rate, and back-to-back frames. RFC 2544 performance measurement testing for Layer 2 E-LAN services on MX104 routers supports UNI-to-UNI unicast traffic only. You can enable reflection at the VPLS user-to-network interface (UNI). The following features are also supported:
  - RFC2544 signature check—Verifies the signature pattern in the RFC2544 packets, by default.
  - MAC swap for pseudowire egress reflection—Swaps the MAC addresses for pseudowire reflection.
  - Ether type filter for both pseudowire and Layer 2 reflection—Specifies the ether type used for reflection.
- **Support for PCP version 2 (MX Series)**—Starting in Release 15.1, Junos OS supports Port Control Protocol (PCP) version 2, defined by IETF RFC 6887. PCP version 2 uses the client once for authentication. Junos OS is able to decode and process version 2 and version 1 messages. There are no CLI changes for PCP version 2 support.  
[See [Port Control Protocol Overview](#).]
- **Support for inline MLPPP interface bundles on Channelized E1/T1 Circuit Emulation MICs (MX80, MX104, MX240, MX480, and MX960)**—Starting in Junos OS Release 15.1, you can configure inline MLPPP interfaces on MX80, MX104, MX240, MX480, and MX960 routers with Channelized E1/T1 Circuit Emulation MICs. Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC). The inline LSQ logical interface (lsq-slot/pic/0) is a virtual service

logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC. A maximum of up to eight inline MLPPP interface bundles are supported on Channelized E1/T1 Circuit Emulation MICs, similar to the support for inline MLPPP bundles on other MICs with which they are compatible.

### Software Defined Networking

- **OpenFlow support (MX2010 and MX2020)**—Starting with Junos OS Release 15.1R2, the MX2010 and MX2020 routers support OpenFlow v1.0 and v1.3.1. OpenFlow enables you to control traffic in an existing network using a remote controller by adding, deleting, and modifying flows on a switch. You can configure one OpenFlow virtual switch and one active OpenFlow controller at the `[edit protocols openflow]` hierarchy level on each device running Junos OS that supports OpenFlow. You can also direct traffic from OpenFlow networks over MPLS networks by using logical tunnel interfaces and MPLS LSP tunnel cross-connects.

[See [Understanding Support for OpenFlow on Devices Running Junos OS.](#)]

- **OVSDB support (MX2010 and MX2020)**—Starting with Junos OS Release 15.1R2, the Junos OS implementation of the Open vSwitch Database (OVSDB) management protocol provides a means through which VMware NSX controllers and MX2010 and MX2020 routers that support OVSDB can communicate.

In an NSX multi-hypervisor environment, NSX controllers and MX2010 and MX2020 routers can exchange control and statistical information via the OVSDB schema, thereby enabling virtual machine (VM) traffic from entities in a virtual network to be forwarded to entities in a physical network and the reverse.

[See [Understanding the OVSDB Protocol Running on Juniper Networks Devices.](#)]

### Software Installation and Upgrade

- **Validate system software add against running configuration on remote host or routing engine**—Beginning with Junos OS Release 15.1R2, you can use the `validate-on-host hostname` and `validate-on-routing-engine routing-engine` options with the `request system software add package-name` command to verify a candidate software bundle against the running configuration on the specified remote host or Routing Engine.
- **Validate system software against running configuration on remote host**—Beginning with Junos OS Release 15.1R2, 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.
- **Support for FreeBSD 10 kernel for Junos OS (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1, on the MX240, MX480, MX960, MX2010, and MX2020 only, FreeBSD 10 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. Because installation restructures the file system, logs and configurations are lost unless precautions are taken. There are now Junos OS and OAM volumes, which provide the ability to boot from the OAM

volume upon failures. Some system commands display different output and a few others are deprecated.

[See [Understanding Junos OS with Upgraded FreeBSD](#).]

---

## Subscriber Management and Services (MX Series)

---



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

- 
- **Additional IPsec encryption algorithms added to support IPsec update data path processing (MX Series)**—Starting in Junos OS Release 15.1, you can configure three new IPsec encryption algorithm options for manual Security Associations at the `[edit security ipsec security-association sa-name manual direction encryption]` hierarchy level: `aes-128-cbc`, `aes-192-cbc`, and `aes-256-cbc`.

[See [encryption \(Junos OS\)](#).]

- **Captive portal content delivery (HTTP redirect) supported on MS-MICs, MS-MPCs, and the Routing Engine (MX Series)**—Starting in Junos OS Release 15.1, you can configure the captive portal content delivery (HTTP redirect) service package for installation using the `set chassis` operational mode command. You can deploy HTTP redirect functionality with a local server or a remote server. The Routing Engine-based captive portal supports a walled garden as a firewall service filter only.

[See [HTTP Redirect Service Overview](#).]

- **LNS support for IPv6-only configurations (MX Series)**—Starting in Junos OS Release 15.1, L2TP LNS supports IPv6-only configurations, in addition to existing IPv4-only and dual-stack configurations. Include the `family inet6` statement in the dynamic profile for IPv6-only dynamic LNS sessions. In earlier releases, LNS supports IPv4-only and dual-stack IPv4/IPv6 configurations.



**NOTE:**

Dynamic LNS sessions require you to include the `dial-options` statement in the dynamic profile, which in turn requires you to include the `family inet` statement. This means that you must include the address families as follows:

- IPv4-only LNS sessions: `family inet`
- IPv6-only LNS sessions: `family inet` and `family inet6`
- Dual-stack IPv4/IPv6 LNS sessions: `family inet` and `family inet6`

---

[See [Configuring a Dynamic Profile for Dynamic LNS Sessions](#).]

- **MAC address option for the Calling-Station-ID attribute (MX Series)**—Starting in Junos OS Release 15.1, you can specify that the subscriber MAC address is included in

the Calling-Station-ID RADIUS attribute (31) that is passed to the RADIUS server. To do so, include the **mac-address** option when you configure the **calling-station-id-format** statement at the **[edit access profile *profile-name* radius options]** hierarchy level.

When all format options are configured, they are ordered in the Calling-Station-Id as follows:

```
nas-identifier#interface description#interface text
description#agent-circuit-id#agent-remote-id#mac address#stacked vlan#vlan
```

[See [Configuring a Calling-Station-ID with Additional Attributes.](#)]

- **Support for overriding L2TP result codes (MX Series)**—Starting in Junos OS Release 15.1, you can configure the LNS to override result codes 4 and 5 with result code 2 in Call-Disconnect-Notify (CDN) messages. These result codes indicate that the number of L2TP sessions have reached the configured maximum value and the LNS can support no more sessions. When the LAC receives the code, it fails over to another LNS to establish subsequent sessions. Some third-party LACs respond only to result code 2.

Include the **override-result-code session-out-of-resource** statement at the **[edit access-profile *access-profile-name* client *client-name* l2tp]** hierarchy level. Issue the **show services l2tp detail | extensive** command to display whether the override is enabled.

[See [override-result-code \(L2TP Profile\).](#)]

- **Support for up to 256 L2TP tunnel groups (MX Series)**—Starting in Junos OS Release 15.1, you can configure and commit up to 256 tunnel groups. In earlier releases, the CLI prevents you from committing the configuration when you create more than 32 groups.

[See [Configuring an L2TP Tunnel Group for LNS Sessions with Inline Services Interfaces.](#)]

- **DHCPv6 relay agent Remote-ID (option 37) based on DHCPv4 relay agent information option 82 (MX Series)**—Starting in Junos OS Release 15.1, DHCPv6 relay agent supports a Remote-ID option (option 37) that is based on the DHCPv4 relay agent information option (option 82). When you enable this feature in dual-stack environments, the DHCPv6 relay agent checks the DHCPv4 binding for the option 82 Remote-ID suboption (suboption 2) and uses that information as option 37 in the outgoing RELAY-FORW message. In addition, you can specify the action DHCPv6 relay agent takes if the DHCPv4 binding does not include an option 82 suboption 2 value; either forward the Solicit message without option 37 or drop the message.

[See [Inserting DHCPv6 Remote-ID Option \(Option 37\) In DHCPv6 Packets.](#)]

- **Juniper Networks VSA 26-181 (DHCPv6-Guided-Relay-Server) support (MX Series)**—Starting in Junos OS Release 15.1, the Junos OS AAA implementation supports Juniper Networks VSA 26-181 (DHCPv6-Guided-Relay-Server). The new support enables RADIUS to use Access-Accept messages to specify the addresses of the DHCPv6 servers to which the DHCPv6 relay agent sends Solicit and subsequent DHCPv6 messages for particular clients. The list of DHCPv6 servers specified by VSA 26-181 takes precedence over the locally configured DHCPv6 server groups for the particular client. You use multiple instances of VSA 26-181 to specify a list of DHCPv6 servers. Creating a list of servers provides load balancing for your DHCPv6 servers, and also enables you to specify explicit servers for a specific client.

[See [Juniper Networks VSAs Supported by the AAA Service Framework.](#)]

- **Asynchronous single hop BFD support for IP liveness detection (MX Series)**—Starting in Junos OS Release 15.1, Bidirectional Forwarding Detection (BFD) supports Layer 3 liveness detection of IP sessions between the broadband network gateway (BNG) and customer premises equipment (CPE). You can show all BFD sessions for subscribers using the **show bfd subscriber session** operational mode command.

[See [show bfd subscriber session](#).]

- **IP session monitoring for DHCP subscribers using the BFD protocol support for active session health checks (MX Series)**—Starting in Junos OS Release 15.1, you can configure a DHCP local server, or DHCP relay agent, or DHCP relay proxy agent to periodically initiate a live detection request to an allocated subscriber IP address of every bound client that is configured to be monitored by using the BFD protocol as the liveness detection mechanism. If a given subscriber fails to respond to a configured number of liveness detection requests, then that subscriber's binding is deleted and its resources released.

[See [DHCP Liveness Detection Overview](#).]

- **IPCP negotiation with optional peer IP address (MX Series)**—Starting in Junos OS Release 15.1, you can configure the **peer-ip-address-optional** statement to enable the Internet Protocol Control Protocol (IPCP) negotiation to succeed even though the peer does not include the IP address option in an IPCP configuration request for static and dynamic, and terminated and tunneled, Point-to-Point Protocol over Ethernet (PPPoE) subscribers. By default, this statement is disabled. This feature also supports high availability (HA) and unified in-service software upgrade (ISSU).

You must assign an IP address by configuring the Framed-IP-Address RADIUS attribute, or the Framed-Pool RADIUS attribute, or by allocating an IP address from the local address pool without a RADIUS-specified pool name, with an optional Framed-Route RADIUS attribute returned from the RADIUS Server.

[See [peer-ip-address-optional](#).]

- **Support for enhanced subscriber management subscriber logical interfaces or interface sets over MPLS pseudowires for a CoS scheduler hierarchy (MX Series)**—Starting in Junos OS Release 15.1, you can enable a CoS scheduling hierarchy for enhanced subscriber management subscriber logical interfaces or interface sets over MPLS pseudowire logical interfaces to function as a Layer 2 node in a CoS hierarchical scheduler. A subscriber logical interface is considered to operate at Layer 2 only if you configure three-level hierarchical scheduling on the logical tunnel anchor point on the physical interface (the IFD). An MPLS pseudowire is a virtual device that is stacked above the logical tunnel anchor point. Implicit hierarchy processes the interface stack properly in such a setup. To configure three-level hierarchical scheduling, include the **implicit-hierarchy** statement at the **[edit interfaces "\$junos-interface-ifd-name" hierarchical-scheduler]** or **[edit interfaces lt-device hierarchical-scheduler]** hierarchy level.
- **Support for enhanced subscriber management subscriber logical interfaces or interface sets over underlying logical interfaces for a CoS scheduler hierarchy (MX Series)**—Starting in Junos OS Release 15.1, you can enable a CoS scheduling hierarchy for enhanced subscriber management subscriber logical interfaces or interface sets over underlying logical interfaces. In Junos OS Release 14.2 and earlier, an interface

set can be either at Layer 2 or Layer 3 levels of the CoS three-level hierarchical scheduler. You can now enable an enhanced subscriber management subscriber logical interface, such as an underlying logical interface, to function as a Layer 2 node in a CoS hierarchical scheduler. To configure three-level hierarchical scheduling, include the **implicit-hierarchy** statement at the **[edit interfaces “\$junos-interface-ifd-name” hierarchical-scheduler]** or **[edit interfaces lt-device hierarchical-scheduler]** hierarchy level.

- **Enhanced subscriber management support for ACI-based PPPoE subscriber session lockout (MX Series)**—Starting in Junos OS Release 15.1, enhanced subscriber management supports identification and filtering of PPPoE subscriber sessions by either the agent circuit identifier (ACI) value or the unique media access control (MAC) source address. You can use this feature when you configure PPPoE subscriber session lockout on static or dynamic VLAN and static or dynamic VLAN demux underlying interfaces.

ACI-based or MAC-based PPPoE subscriber session lockout prevents a failed or short-lived PPPoE subscriber session from reconnecting to the router for a default or configurable time period. ACI-based PPPoE subscriber session lockout is useful for configurations such as PPPoE interworking in which MAC source addresses are not unique on the PPPoE underlying interface.

To configure ACI-based PPPoE subscriber session lockout for enhanced subscriber management, use the same procedure that you use to configure it on a router without enhanced subscriber management enabled.

[See [PPPoE Subscriber Session Lockout Overview](#).]

- **Subscriber Secure Policy (SSP) interception of Layer 2 datagrams (MX Series)**—Starting in Junos OS Release 15.1, when DTCP- or RADIUS-initiated SSP intercepts traffic on a logical subscriber interface, including VLAN interfaces, the software intercepts Layer 2 datagrams and sends them to the mediation device. Previously, the software intercepted Layer 3 datagrams on logical subscriber interfaces.

Interception of subscriber traffic on an L2TP LAC interface is unchanged. The Junos OS software sends the entire HDLC frame to the mediation device.

Interception of subscriber traffic based on interface family, such as IPv4 or IPv6, is also unchanged. The Junos OS software sends the Layer 3 datagram to the mediation device.

Interception of traffic based on a subscriber joining a multicast group is also unchanged. Layer 3 multicast traffic is intercepted and sent to the mediation device. However, multicast traffic that passes through a logical subscriber interface is intercepted along with other subscriber traffic, and is sent as a Layer 2 datagram to the mediation device.

[See [Subscriber Secure Policy Overview](#).]

- **Additional methods to derive values for L2TP connect speeds (MX Series)**—Starting in Junos OS Release 15.1, several new ways are supported for determining the transmit and receive connect speeds that the LAC sends to the LNS:
  - The Juniper Networks VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163), can provide the values.

- The Juniper Networks VSA, Tunnel-Tx-Speed-Method (26-94), can specify a method (source) for the LAC to derive the values.
- You can configure the LAC to use the actual downstream traffic rate enforced by CoS for the transmit speed. The **actual** method requires the effective shaping rate to be enabled and does not provide a receive speed, which is determined by the fallback scheme.

You can also configure the LAC not to send the connect speeds.

[See [Transmission of Tx Connect-Speed and Rx Connect-Speed AVPs from LAC to LNS.](#)]

- **Pseudowire device support for reverse-path forwarding check (MX Series)**—Starting in Junos OS Release 15.1, unicast reverse-path forwarding checks are supported on pseudowire subscriber logical interface devices (ps0) for both the inet and inet6 address families. Include the **rpf-check** statement at the **[edit interfaces ps0 unit logical-unit-number family family]** hierarchy level for either address family.

[See [Configuring a Pseudowire Subscriber Logical Interface Device.](#)]

- **Destination-equal load balancing for L2TP sessions (MX Series)**—Starting in Junos OS Release 15.1, you can enable the LAC to balance the L2TP session load equally across all tunnels at the highest available preference level by evaluating the number of sessions to the destinations and the number of sessions carried by the tunnels. By default, tunnel selection within a preference level is strictly random. Include the **destination-equal-load-balancing** statement at the **[edit services l2tp]** hierarchy level to load-balance the sessions. The **weighted-load-balancing** statement must be disabled.

[See [LAC Tunnel Selection Overview](#) and [Configuring Destination-Equal Load Balancing for LAC Tunnel Sessions.](#)]

- **Support for Extensible Subscriber Services Manager (MX Series)**—Starting in Release 15.1, Junos OS supports Extensible Subscriber Services Manager (ESSM), a background process that enables dynamic provisioning of business services.
- **Loopback address as source address on DHCP relay agent**—Starting in Junos OS Release 15.1, you can configure the DHCPv4 and DHCPv6 relay agent to use the relay agent loopback address as the source address in DHCP packets. In network configurations where a firewall on the Border Network Gateway (BNG) is between the DHCP relay agent and the DHCP server, the BNG firewall recognizes only the BNG loopback address. In that case, DHCP unicast packets are not recognized. You can use two new configuration statements to override the DHCP source address with the BNG loopback address so DHCP packets are recognized.

For both DHCPv4 and DHCPv6, use the **relay-source lo0** statement at the **[edit forwarding-options dhcp-relay group group-name overrides]** hierarchy level to set the source address to the loopback address in DHCP packets.

For DHCPv4, to add the options **link-selection** and **server-id override** to option-82 packets relayed to the DHCP server, use the **server-id-override** statement at one of the following hierarchy levels, either globally or within an interface group:

- **[edit forwarding-options dhcp-relay relay-option-82]**
- **[edit forwarding-options dhcp-relay group group-name relay-option-82]**



- [edit logical-systems *logical-system-name* forwarding-options dhcp-relay relay-option-82]
- [edit logical-systems *logical-system-name* forwarding-options dhcp-relay group *group-name* relay-option-82]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-82]
- [edit logical-systems *logical-system-name* routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* relay-option-82]
- [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay relay-option-82]
- [edit routing-instances *routing-instance-name* forwarding-options dhcp-relay group *group-name* relay-option-82]
- **Support for DUID based on link-layer address in DHCPv6**—Starting in Junos OS Release 15.1, the DHCPv6 server supports clients using a DHCP Unique ID (DUID) based on link-layer address (DUID-LL). To change from the default vendor-assigned DUID based on enterprise number (DUID-EN) to DUID-LL, use the new **server-duid-type duid-ll** configuration statement at the [edit system services dhcp-local-server dhcpv6] hierarchy level.
- **New support for Framed-IP-Netmask for access-internal routes (MX Series)**—Starting in Junos OS Release 15.1R2, the mask value returned by RADIUS in the Framed-IP-Netmask attribute during PPP negotiation is considered for application to the access-internal route for the subscriber session. In earlier releases, the attribute mask is ignored and a /32 netmask is always applied, with the consequence that the address is set to the value of the Framed-IP-Address attribute returned by RADIUS.

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

When the SDB\_FRAMED\_PROTOCOL attribute is equal to AUTHD\_FRAMED\_PROTOCOL\_PPP, the **show subscribers** command now displays the actual value of Framed-IP-Netmask in the IP Netmask field. Otherwise, the field displays the default value of 255.255.255.255.

- **Support for saving accounting files when Routing Engine mastership changes (MX Series)**—Starting in Junos OS Release 15.1R2, you can configure the router to save the accounting files from the new backup Routing Engine to the new master Routing Engine when a change in mastership occurs. To do so, include the **push-backup-to-master** statement at the [edit accounting-options file *filename*] hierarchy level.

Configure this option when the new backup Routing Engine is not able to connect to the archive site; for example, when the site is not connected by means of an out-of-band interface or the path to the site is routed through a line card. The files are stored in the **/var/log/pfedBackup** directory on the router. The master Routing Engine includes these accounting files with its own current accounting files when it transfers the files from the backup directory to the archive site at the next transfer interval.

- **Disabling DHCP snooping filters for DHCP traffic that can be directly forwarded (MX Series)**—Starting in Junos OS Release 15.1R2, you can disable DHCP snooping filters for an address family in the routing context in which snooping is configured.

When you first enable DHCP snooping, all DHCP traffic is snooped by default and only DHCP packets associated with subscribers (or their creation) will be handled, all other DHCP packets will be discarded. You can optionally modify this dropping behavior based on the type of interface—configured interfaces, non-configured interfaces, or all interfaces. All snooped DHCP traffic is still forwarded to the routing plane in the routing instance, and in some cases, this results in excessive DHCP traffic being sent to the routing plane for snooping. The **no-snoop** statement disables snooping filters for DHCP traffic that can be forwarded directly from the hardware control plane, such as Layer 3 unicast traffic with a valid route, preventing that DHCP traffic from being forwarded to the slower routing plane of the routing instance.

- **Disabling DHCP snooping filters for DHCP traffic that can be directly forwarded (MX Series)**—Starting in Junos OS Release 15.1R2, you can disable DHCP snooping filters for an address family in the routing context in which snooping is configured.

When you first enable DHCP snooping, all DHCP traffic is snooped by default and only DHCP packets associated with subscribers (or their creation) will be handled, all other DHCP packets will be discarded. You can optionally modify this dropping behavior based on the type of interface—configured interfaces, non-configured interfaces, or all interfaces. All snooped DHCP traffic is still forwarded to the routing plane in the routing instance, and in some cases, this results in excessive DHCP traffic being sent to the routing plane for snooping. The **no-snoop** statement disables snooping filters for DHCP traffic that can be forwarded directly from the hardware control plane, such as Layer 3 unicast traffic with a valid route, preventing that DHCP traffic from being forwarded to the slower routing plane of the routing instance.

---

## User Interface and Configuration

- **Support for displaying configuration differences in XML tag format (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, you can use the **show compare | display xml** command to compare the candidate configuration with the current committed configuration and display the differences between the two configurations in XML tag format.

[See [Understanding the show | compare | display xml Command Output](#).]

---

## VPNs

- **Leveraging DPCs for EVPN deployment (MX Series routers with DPCs)**—Starting with Junos OS Release 15.1, DPCs can be leveraged to provide support for Ethernet VPN (EVPN) functionality. Earlier, the EVPN functionality was provided by MX Series routers with MPC and MIC interfaces only.

The DPC support for EVPN is provided with the following considerations:

- DPCs provide support for EVPN in the active/standby mode of operation including support for the following:
  - EVPN instance (EVI)

- Virtual switch (VS)
- Integrated routing and bridging (IRB) interfaces
- DPCs intended for providing the EVPN active/standby mode support should be the customer edge (CE) device-facing line card. The provider edge (PE) device interfaces in the EVPN domain should use only MPC and MIC interfaces.

[See [EVPN Multihoming Overview](#).]



**NOTE:** Although present in the code, the Ethernet VPN (EVPN) active/active multihoming feature is not supported in Junos OS Release 15.1R2.

**Active/active multihoming support for EVPNs (MX Series routers with MPCs and MICs only)**—The Ethernet VPN (EVPN) solution on MX Series routers with MPC and MIC interfaces is extended to provide multihoming functionality in the active/active redundancy mode of operation. This feature enables load-balancing of Layer 2 unicast traffic across all the multihomed links on and toward a customer edge device, and provides link-level and node-level redundancy along with effective utilization of resources.

- **Enhanced Group VPNv2 member features (MX10, MX20, MX40, MX80, MX240, MX480, MX960)**—Starting in Junos OS Release 15.1, Group VPNv2 member features have been enhanced to include the following:
  - Accept group domain of interpretation (GDOI) push messages from Cisco group controller/key server (GC/KS) as per RFC 6407.
  - Support for group associated policy (GAP) payload, including activation time delay (ATD) and deactivation time delay (DTD), in push messages from Cisco GC/KS as per RFC 6407.
  - Support standardized push ACK messages from MX Series group member router to Cisco GC/KS as per IETF draft RFC <http://www.ietf.org/id/draft-weis-gdoi-rekey-ack-00.txt>.
  - IP Delivery Delayed Detection Protocol. Time-based anti-replay protection for Group VPNv2 data traffic on MX Series group member routers as per IETF draft RFC <http://tools.ietf.org/html/draft-weis-delay-detection-00>.
  - Support for SHA-256 HMAC algorithm for authentication.
  - Support partial fail open for business-critical traffic.
  - Support for control-plane debug traces per member IP address and server IP address.
  - Same gateway for multiple groups, wherein the same local and remote address pair is used for multiple groups.

[See [Group VPNv2 Overview](#).]

- **Segmented inter-area P2MP LSP (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, P2MP LSPs can be segmented at the area boundary. A

segmented P2MP LSP consists of an ingress area segment (ingress PE router or ASBR), backbone area segment (Transit ABR), and egress area segment (egress PE routers or ASBRs). Each of the intra-area segments can be carried over provider tunnels such as P2MP RSVP-TE LSP, P2MP mLDP LSP, or ingress replication. Segmentation of inter-area P2MP LSP occurs when the S-PMSI auto-discovery routes are advertised, which triggers the inclusion of a new BGP extended community or inter-area P2MP segmented next-hop extended community in the ingress PE router or ASBR, transit ABR, and egress PE routers or ASBRs.

[See [Example: Configuring Segmented Inter-Area P2MP LSP](#).]

**Related  
Documentation**

- [Changes in Behavior and Syntax on page 52](#)
- [Known Behavior on page 67](#)
- [Known Issues on page 68](#)
- [Resolved Issues on page 76](#)
- [Documentation Updates on page 99](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)
- [Product Compatibility on page 113](#)

## 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 15.1R2 for the M Series, MX Series, and T Series.

- [Authentication, Authorization and Accounting on page 53](#)
- [Class of Service \(CoS\) on page 53](#)
- [General Routing on page 53](#)
- [High Availability \(HA\) and Resiliency on page 53](#)
- [Junos XML API and Scripting on page 55](#)
- [Layer 2 VPNs on page 55](#)
- [MPLS on page 55](#)
- [Multicast on page 55](#)
- [Network Management and Monitoring on page 55](#)
- [Routing Policy and Firewall Filters on page 56](#)
- [Routing Protocols on page 56](#)
- [Security on page 59](#)
- [Services Applications on page 59](#)
- [Subscriber Management and Services \(MX Series\) on page 60](#)
- [System Logging on page 65](#)
- [System Management on page 66](#)

- [User Interface and Configuration on page 66](#)
- [VPNs on page 66](#)

---

### Authentication, Authorization and Accounting

- **Statement introduced to enforce strict authorization**—Starting in Junos OS Release 15.1R2, customers can use the **set system tacplus-options strict-authorization** statement to enforce strict authorization to the users. When a user is logging in, Junos OS issues two TACACS+ requests—first is the authentication request and then the authorization request. By default, when the authorization request is rejected by the TACACS+ server, Junos OS ignores this and allows full access to the user. When the **set system tacplus-options strict-authorization** statement is set, Junos OS denies access to the user even on failure of the authorization request.

---

### Class of Service (CoS)

- **Change to CoS shaping rate fallback behavior (MX Series)**—Starting in Junos OS Release 15.1, when a CoS service profile is deactivated, the traffic shaping rate falls back in the following order: ANCP shaping rate, PPPoE IA tag rate, or shaping rate configured in the traffic control profile. In earlier releases, the traffic shaping rate falls back to the ANCP adjusted rate or the traffic control profile value.

Now when an ANCP shaping rate adjustment is removed, the rate falls back to the PPPoE IA tag rate or the traffic control profile value. In earlier releases, the rate falls back to the traffic control profile value.

[See [CoS Adjustment Control Profiles Overview](#).]

---

### General Routing

- **commit synchronize statement is not allowed in batch mode**—When you attempt to execute **commit atomic** in configure batch mode, a warning message is displayed: **warning: graceful-switchover is enabled, commit synchronize should be used**. This is because commit synchronize is not allowed to be given in configure batch mode. In this case, issue the **set system commit synchronize** command followed by **commit**.

---

### High Availability (HA) and Resiliency

- **VRRP adjusted priority can go to zero (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, the adjusted priority of a configured VRRP group can go to zero (0). A zero (0) priority value is used to trigger one of the backup routers in a VRRP group to quickly transition to the master router without having to wait for the current master to timeout. Prior to Junos OS Release 15.1, an adjusted priority could not be zero. This change in behavior prevents the VRRP group from blackholing traffic.

[See [Configuring a Logical Interface to Be Tracked for a VRRP Group](#) or [Configuring a Route to Be Tracked for a VRRP Group](#).]

- **A check option is added for command request chassis routing-engine master**—Starting in Junos OS Release 15.1, a **check** option available with the **switch**, **release**, and **acquire** options checks the GRES status of the standby Routing Engine before toggling mastership. The **force** option is also removed.

[See [request chassis routing-engine master](#).]

- **GRES readiness is part of show system switchover output (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, switchover readiness status is reported as part of the output for the operational mode command **show system switchover**. This is true for the TX Matrix Plus platform as well.

[See [show system switchover](#).]

- **Improved command output for determining GRES readiness in an MX Series Virtual Chassis (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, the **request virtual-chassis routing-engine master switch check** command displays the following output when the member routers in a Virtual Chassis are ready to perform a graceful Routing Engine switchover (GRES):

```
{master:member0-re0}
```

```
user@host> request virtual-chassis routing-engine master switch check
Switchover Ready
```

In earlier releases, the **request virtual-chassis routing-engine master switch check** command displays no output to confirm that the member routers are ready for GRES.

The output of the **request virtual-chassis routing-engine master switch check** command has not changed when the member routers are not yet ready for GRES.

[See [Determining GRES Readiness in a Virtual Chassis Configuration](#).]



**NOTE:** The changes to global switchover behavior in an MX Series Virtual Chassis are *not supported* in Junos OS Release 15.1. Documentation for this feature is included in the Junos OS 15.1 documentation set.

**Changes to global switchover behavior in an MX Series Virtual Chassis (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, performing a global switchover by issuing the **request virtual-chassis routing-engine master switch** command from the master Routing Engine in the Virtual Chassis master router (VC-M) has the same result as performing a local switchover from the VC-M.

After a global switchover, the Virtual Chassis master router (VC-M) becomes the Virtual Chassis backup router (VC-B), and the VC-B becomes the VC-M. In addition, a global switchover now causes the local roles (**master** and **standby**, or **m** and **s**) of the Routing Engines in the former VC-M to change, but does not change the local roles of the Routing Engines in the former VC-B.

In earlier releases, a global switchover in a Virtual Chassis caused the VC-M and VC-B to switch global roles, but did not change the master and standby local roles of the Routing Engines in either member of the Virtual Chassis.

[See [Switchover Behavior in an MX Series Virtual Chassis](#).]

- **New unified 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” to confirm whether you want to proceed with the ISSU operation or not.

### Junos XML API and Scripting

---

- **Escaping of special XML characters required for request\_login (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 15.1R2, you must escape any special characters in the username and password elements of a **request\_login** XML RPC request. The following five symbols are considered special characters: greater than (>), less than (<), single quote ('), double quote ("), and ampersand (&). Both entity references and character references are acceptable escape sequence formats. For example, **&amp;** and **&#38;** are valid representations of an ampersand. Previously no escaping of these characters was required.

### Layer 2 VPNs

---

- **Support for hot standby pseudowire for VPLS instances with LDP (MX Series)**—Starting with Junos OS Release 15.1R2, you can configure a routing device running a VPLS routing instance configured with the Label Distribution Protocol (LDP) to indicate that a hot-standby pseudowire is desired upon arrival of a PW\_FWD\_STDBY status-tlv. Include the **hot-standby-vc-on** statement at the **[edit routing instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor address pseudowire-status-tlv]** hierarchy level.

### MPLS

---

- **Deselecting active path on bandwidth reservation failure (MX Series)**—LSP deselects the current active path if the path is not able to reserve the required amount of bandwidth and there is another path that is successful and capable of becoming active. If the current active path is not deselected, then it continues to be active despite having insufficient bandwidth. If none of the paths are able to reserve the required amount of bandwidth, then the **tear-lsp** option brings down the LSP.

[See [deselect-on-bandwidth-failure](#).]

### Multicast

---

- **Disabling igmp-snooping on VPLS (MX Series)**—In order to make configuration and debugging easier, starting in Junos OS Release 15.1, multiple Group VPNv2 groups can use the same gateway. The commit check for a unique tuple of <**local\_address**, **remote\_address**, **routing\_instance**> across groups has been removed. The same tuple is now checked for uniqueness across all gateways. This allows multiple groups to share the same gateway for their Group VPNv2 traffic.

### Network Management and Monitoring

---

- **Enhanced service type information in an SNMP MIB walk operation for jnxSpSvcSet**—Starting with releases 13.3R7, 14.1R6, 14.2R4, and 15.1R2, Junos OS provides enhanced service type (SvcType) information in a MIB walk operation for the jnxSpSvcSet MIB table. Stateful firewall, NAT, and IDS service sets are now categorized

under the **SFW/NAT/IDS** service type. IPsec services are categorized as **IPSEC** service type, while all other services are grouped as **EXT-PKG**.

In Junos OS Release 13.3R6 and earlier, the **show snmp mib walk** command for the `jnxSpSvcSet` MIB table displays the service type as **EXT-PKG** for all services.

- **SNMP proxy feature (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1R2, you must configure the **interface <interface-name>** statement at the **[edit snmp]** hierarchy level for the proxy SNMP agent. Earlier, configuring an interface for the proxy SNMP agent was not mandatory.
- **Change in how used memory is calculated in Junos OS with upgraded FreeBSD (MX Series)**—Starting in Junos OS Release 15.1, for platforms running Junos OS with upgraded FreeBSD, the way used memory is calculated has changed. Inactive memory is no longer included in the calculation for memory utilization. This change is reflected in the value given for memory utilization in the output for the **show chassis routing-engine** command. This change also affects the SNMP representation of this value at `jnxOperatingBuffer`.

[For platforms that run Junos OS with upgraded FreeBSD, see [Understanding Junos OS with Upgraded FreeBSD](#).]

---

## Routing Policy and Firewall Filters

- **Command completion for the [show firewall prefix-action-stats filter *filter-name* prefix-action] hierarchy on all compatible platforms**—In releases earlier than Junos OS Release 15.1, you could not utilize the command completion feature at the **[show firewall prefix-action-stats filter *filter-name* prefix-action]** hierarchy level. This meant that you had to know the name of the prefix-action in order to complete any command at that hierarchy level. This involved running a show configuration command, getting the prefix-action name, and using it in the command.

Starting in Junos OS Release 15.1, command completion is available so that pressing the Tab key at the **[show firewall prefix-action-stats filter *filter-name* prefix-action]** hierarchy level lists all currently configured prefix-action names.

---

## Routing Protocols

- **Enhanced show isis overview command (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 15.1, the **show isis overview** command display output includes details, such as **Hostname**, **Sysid**, and **Areaid**. This additional information facilitates troubleshooting IS-IS adjacency issues.

[See [show isis overview](#).]

- **RPD refreshes the route record database only if there is a new update (MX Series)**—Beginning with Junos OS Release 15.1, when you commit a minor configuration change, the rpd sends only AS paths that are active routes to the FPCs. Not all known AS paths are sent to the FPC, thereby considerably reducing the memory and CPU usage, resulting in a faster route record database update. Route record now keeps track of configuration and reconfiguration times. At client startup, all the routes are sent to the client, but at reconfiguration, route record now checks the timestamp of the route.



In earlier Junos OS releases, when a configuration change was committed, the Routing Engine CPU usage and the FPC CPU usage would go high for an extended period of time. This occurred even if there was a minor change to the configuration. The FPCs and the client were running out of memory due to the high number of AS paths sent by route record. This was especially evident in very large-scale configurations where the number of AS paths and the number of routes were large. This took a lot of CPU time and memory to process because at reconfiguration, route record sent all routes to the client again, even if there were no route changes.

- **BGP hides a route received with a label block size greater than 256 (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, when a BGP peer (running Junos OS) sends a route with a label block size greater than 256, the local speaker hides the route and does not re-advertise this route. The output of the **show route detail/extensive hidden/all** displays the hidden route and states the reason as **label block size exceeds max supported value**. In earlier Junos OS releases, when a peer sent a route with a label block size greater than 256, the routing protocol process (rpd) terminated abnormally.
- **New option to remove peer loop check (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, a new option **no-peer-loop-check** to remove the peer loop check for private AS numbers is available under the **remove-private** statement at the following hierarchy levels:
 

```
[edit logical-systems logical-system-name protocols bgp]
[edit protocols bgp]
[edit routing-instances routing-instance-name protocols bgp]
```
- **BGP link state value modified to 29 (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.2R3, the value of the BGP **LINK-STATE** (LS) path attribute is modified to 29, which is IANA's officially assigned value. In earlier Junos OS releases, the **LINK-STATE** path attribute had a private value of 99 that was used for interoperability testing with other vendors. The previous versions of BGP LS are not compatible with this new value of BGP LS. Therefore, BGP LS users cannot use unified ISSU with the BGP LS value of 29.
- **DSCP bit not copied into IPv6 ICMP reply packets (MX Series)**—Beginning with Junos OS Release 15.1, the Differentiated Services code point (DSCP) field from the IPv6 header of the incoming ICMP request packet is copied into the ICMP reply packet. The value of the DSCP field represents the class of service, and transmission of packets is prioritized based on this value. In earlier Junos OS releases, the value of the DSCP field was set to 0, which is undesirable because the class of service information is lost. Junos OS now retains the value of the DSCP field in the incoming packet and copies it into the ICMP reply packet.
- **New IS-IS adjacency holddown CLI command (MX Series)**—Beginning with Junos OS Release 15.1, a new operational command **show isis adjacency holddown** is introduced to display the adjacency holddown status. This command is useful to verify whether the adjacency holddown is enabled and facilitates troubleshooting when there are adjacency issues due to IS-IS adjacency holddown.

[See [show isis adjacency holddown](#).]

- **Eliminate fe80::/64 direct routes from RIB for IPv6 interfaces**—Beginning with Junos OS Release 15.1, the fe80::/64 direct routes for IPv6 addresses are not installed in the routing table. Therefore, when you issue a **show route** command, the fe80::/64 routes for IPv6 addresses are not displayed in the output. In earlier releases, Junos OS added the fe80::/64 direct routes to the routing table when inet6 family was enabled on an interface. These fe80::/64 direct routes are neither routable nor used for routing decisions and hence their absence in the routing table does not impact any functionality.
- **Configure and establish targeted sessions with third-party controllers using LDP targeted neighbor (M Series and MX Series)**— Starting with Junos OS Release 15.1, you can configure LDP targeted neighbor to third-party controllers for applications such as route recorder that wants to learn label-FEC bindings of an LSR. LDP targeted neighbor helps to establish a targeted session with controllers for a variety of applications.
- **Enable forwarding IPv6 solicited router advertisements as unicast**—Beginning with Junos OS Release 15.1, you can configure devices to send router advertisements as unicast in response to the router solicitation message sent by IPv6 routers. In earlier Junos OS releases, IPv6 router advertisements were sent as periodic multicast, which caused a battery drain in all the other devices. A new configuration statement **solicit-router-advertisement-unicast** is introduced at the **[edit protocols router-advertisement interface *interface-name*]** hierarchy level.

[See [solicit-router-advertisement-unicast](#).]

- **Enhanced BGP log message when prefix limit is exceeded**—Beginning with Junos OS Release 13.3, BGP generates an enhanced log message when the prefix limit exceeds the configured limit. The log message now includes the instance name in addition to the peer address and address family.

[See [prefix-limit](#).]

## Security

- **Packet types added for DDoS protection L2TP policers (MX Series routers with MPCs, T4000 routers with FPC5)**—Starting in Junos OS Release 15.1, the following eight packet types have been added to the DDoS protection L2TP protocol group to provide flexibility in controlling L2TP packets:

cdn	scccn
hello	sccrq
iccn	stopccn
icrq	unclassified

Previously, no individual packet types were available for this protocol group and all L2TP packets were policed the same based on the aggregate policer value. The default values for the bandwidth and burst policers for all packet types is 20,000 pps. The default **recover-time** is 300 seconds for each of the L2TP packet types.

[See [protocols \(DDoS\)](#).]

## Services Applications

- **Support for configuring TWAMP servers on routing instances (MX Series)**—Starting in Junos OS Release 15.1, you can specify the TWAMP servers on specific routing instances, instead of associating the TWAMP server at the system level. To apply the TWAMP server to a routing instance configured on a router, include the **routing-instance-list *instance-name* port *port-number*** statement at the **[edit services rpm twamp server]** hierarchy level. The port number of the specified routing instance is used for TWAMP probes that are received by a TWAMP server. The default routing instance is Internet routing table inet.0. If you do not specify a routing instance, the TWAMP probe applies to all routing instances. To apply the TWAMP probe to only the default routing instance, you must explicitly set the value of instance-name to the default. If an interface is not part of any routing instance, the default port is used for TWAMP probes. You can configure up to 100 routing instances for a TWAMP server.
- **Optional inclusion of Flags field in DTCP LIST messages (MX Series)**—Starting in Junos OS Release 15.1, the Flags field is not a required parameter in the DTCP LIST message. The LIST request is not rejected if the LIST message does not contain the Flags field. If the DTCP LIST message contains the Flags field, the value of that field is processed. If the LIST message does not contain the Flags field, the CRITERIA field parameter is used for the Flags field.
- **Change in support for service options configuration on service PICs at the MS and AMS interface levels (MX Series)**—Starting in Junos OS Release 15.1, when a multiservices PIC (**ms-** interface) is a member interface of an AMS bundle, you can configure the service options to be applied on the interface only at the **ms-** interface level or the AMS bundle level by including the **services-options** statement at the **[edit interfaces *interface-name*]** hierarchy level at a point in time. You cannot define service

options for a service PIC at both the AMS bundle level and at the **ms-** interface level simultaneously. When you define the service options at the MS level or the AMS bundle level, the service options are applied to all the service-sets, on the **ms-** interface or the AMS interface defined at **ms-fpc/pic/port.logical-unit** or **amsN**, respectively.

- **Changes in the format of session open and close system log messages (MX Series routers with MS-MICs and MS-MPCs)**—Starting with Junos OS Release 15.1, with the Junos OS Extension-Provider packages installed and configured on the device for MS-MPCs and MS-MICs, the formats of the **MSVCS\_LOG\_SESSION\_OPEN** and **MSVCS\_LOG\_SESSION\_CLOSE** system log messages are modified to toggle the order of the destination IPv4 address and destination port address displayed in the log messages to be consistent and uniform with the formats of the session open and close logs of MS-DPCs.
- **Support for bouncing service sets for dynamic NAT (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1, for service sets associated with aggregated multiservices (AMS) interfaces, you can configure the **enable-change-on-ams-redistribution** statement at the **[edit services service-set service-set-name service-set-options]** hierarchy level to enable the service set to be bounced (reset) for dynamic NAT scenarios (dynamic NAT, NAT64, and NAT44) when a member interface of an AMS bundle rejoins or a member interface failure occurs. When a member interface fails, the application resources (NAT pool in the case of dynamic NAT scenarios) and traffic load need to be rebalanced. For application resources to be rebalanced, which is the NAT pool for dynamic NAT environments, the NAT pool is split and allocated by the service PIC daemon (spd).
- **Changed range for maximum lifetime for PCP mapping**—Starting in Junos OS Release 15.1, the range for the maximum lifetime, in seconds, for PCP mapping that you can configure by using the **mapping-lifetime-max** statement at the **[edit services pcp]** hierarchy level is modified to be from 0 through 4294667, instead of the previous range from 0 through 2147483647.
- **Change in the test-interval range for RPM tests (MX Series)**—Starting in Junos OS Release 15.1R2, the minimum period for which the RPM client waits between two tests (configured by using the **test-interval** statement at the **[edit services rpm probe owner test test-name]** hierarchy level is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 0 seconds. A test interval of 0 seconds causes the RPM test to stop after one iteration.

---

### Subscriber Management and Services (MX Series)

---



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

---

- **Support for specifying preauthentication port and password (MX Series)**—Starting in Junos OS Release 15.1, you can configure a router that operates as the RADIUS client to contact a RADIUS server for authentication and preauthentication requests on two different UDP ports and using different secret passwords. Similar to configuring the port numbers for authentication and accounting requests, you can define a unique port number that the router uses to contact the RADIUS server for logical line identification (LLID) preauthentication requests. You can also define a unique password for preauthentication requests. If you do not configure a separate UDP port or secret for preauthentication purposes, the same UDP port and secret that you configure for authentication messages is used.

To configure a unique UDP port number and the password to be used to contact the RADIUS server for pre-authentication requests, include the **preauthentication-port** *port-number* and **preauthentication-secret** *password* statements, respectively, at the **[edit access radius-server *server-address*]** or **[edit access profile *profile-name* radius-server *server-address*]** hierarchy level.

[See [Configuring a Port and Password for LLID Preauthentication Requests.](#)]

- **Addition of pw-width option to the nas-port-extended-format statement (MX Series)**—Starting in Junos OS Release 15.1, you can configure the number of bits for the pseudowire field in the extended-format NAS-Port attribute for Ethernet subscribers. Specify the value with the **pw-width** option in the **nas-port-extended-format** statement at the **[edit access profile *profile-name* radius options]** hierarchy level. The configured fields appear in the following order in the binary representation of the extended format:

*aggregated-ethernet slot adapter port pseudo-wire stacked-vlan vlan*

The width value also appears in the Cisco NAS-Port-Info AVP (100). In addition to Junos OS Release 15.1, the **pw-width** option is available in Junos OS Release 13.3R4; it is not available in Junos OS Release 14.1 or Junos OS Release 14.2.

[See [CoS Adjustment Control Profiles Overview.](#)]

- **Enhanced support for Calling-Station-ID (RADIUS attribute 31) (MX Series)**—Starting in Junos OS Release 15.1, you can specify optional information that is included in the Calling-Station-ID that is passed to the RADIUS server. You can now include the following additional information when configuring the **calling-station-id-format** statement at the **[edit access profile *profile-name* radius options]** hierarchy level:
  - **interface-text-description**—Interface description text string
  - **stacked-vlan**—Stacked VLAN ID
  - **vlan**—VLAN ID

[See [Configuring a Calling-Station-ID with Additional Attributes.](#)]

- **Unique RADIUS NAS-Port attributes (MX Series)**—Starting in Junos OS Release 15.1, you can configure unique values for the RADIUS NAS-Port attribute (attribute 5), to ensure that a single NAS-Port attribute is not used by multiple subscribers in the network. You can create NAS-Port values that are unique within the router only, or that are unique across all MX Series routers in the network. To create unique NAS-Port attributes for subscribers, the router uses an internally generated number and an optional unique chassis ID, which you specify. The generated number portion of the

NAS-Port provides uniqueness within the router only. The addition of the optional chassis ID configuration ensures that the NAS-Port is unique across all MX Series routers in the network.

[See [Enabling Unique NAS-Port Attributes \(RADIUS Attribute 5\) for Subscribers.](#)]

- **RADIUS VSA support for IANA Private Enterprise Number 311 primary and secondary DNS servers (MX Series)**—Starting in Junos OS Release 15.1, the Junos OS AAA implementation supports RADIUS VSAs that identify the primary and secondary DNS servers for IANA private enterprise number 311 (Microsoft Corporation). The two VSAs are shown in the following list, and are described in RFC 2548, *Microsoft Vendor-specific RADIUS Attributes*:
  - MS-Primary-DNS-Server (VSA 26-28)—The 4-octet address of the primary Domain Name Server. This VSA can be included in Access-Accept and Accounting-Request packets.
  - MS-Secondary-DNS-Server ((VSA 26-29)—The 4-octet address of the secondary Domain Name Server. This VSA can be included in Access-Accept and Accounting-Request packets.

[See [RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses.](#)]

- **Filters for duplicate RADIUS accounting interim reports (MX Series)**—Starting in Junos OS Release 15.1, subscriber management provides a duplication filter feature that enables you to specify which accounting servers receive RADIUS accounting interim reports when RADIUS accounting duplicate reporting is active. You configure the filters in the AAA access profile, and the router then applies the filters to subscribers associated with that profile.

Subscriber management supports the following filtering for RADIUS accounting duplicate reporting:

- Duplicated accounting interim messages
- Original accounting interim messages
- Excluded RADIUS attributes

Subscriber management also provides additional attribute support for the **exclude** statement at the **[edit access profile *profile-name* radius attributes]** hierarchy level.

[See [Configuring Duplication Filters for RADIUS Accounting Duplicate Reporting.](#)]

- **LAC configuration no longer required for L2TP tunnel switching with RADIUS attributes (MX Series)**—Starting in Junos OS Release 15.1, when you use Juniper Networks VSA 26-91 to provide tunnel profile information for L2TP tunnel switching, you no longer have to configure a tunnel profile on the LAC. In earlier releases, tunnel switching failed when you did not also configure the LAC, even when the RADIUS attributes were present.

[See [Configuring L2TP Tunnel Switching](#) and [L2TP Tunnel Switching Overview.](#)]

- **Changes to ANCP triggering of RADIUS immediate interim accounting updates (MX Series)**—Starting in Junos OS Release 15.1, the AAA daemon immediately sends a RADIUS interim-accounting request to the RADIUS server when it receives notification

of ANCP actual downstream or upstream data rate changes, even when the **update-interval** statement is not included in the subscriber session access profile. In earlier releases, the **update-interval** statement is required. This feature still requires that the **ancp-speed-change-immediate-update** statement is included in the access profile.

[See [Configuring Immediate Interim Accounting Updates to RADIUS in Response to ANCP Notifications](#).]

- **DHCP behavior when renegotiating while in bound state (MX Series)**—Starting in Junos OS Release 15.1, DHCPv4 and DHCPv6 local server and relay agent all use the same default behavior when receiving a DHCPv4 Discover or DHCPv6 Solicit message with a matching client ID, while in a bound state. In the default behavior, DHCP maintains the existing client entry when receiving a new Discover or Solicit message that has a client ID that matches the existing client. In Junos OS releases prior to 15.1, DHCPv6 local server and DHCPv6 relay agent use the opposite default behavior, and tear down the existing client entry when receiving a Solicit message with a matching client ID, while in a bound state.

You use the **delete-binding-on-renegotiation** statement to override the default behavior and configure DHCP local server and relay agent to delete the existing client entry when receiving a Discover or Solicit message while in a bound state.

[See [DHCP Behavior When Renegotiating While in Bound State](#).]

- **Optional CHAP-Challenge attribute configuration (MX Series)**—Starting in Junos OS Release 15.1, you can configure the router to override the default behavior and insert the random challenge generated by the NAS into the Request Authenticator field of Access-Request packets. In the default behavior, the **authd** process sends the random challenge as the CHAP-Challenge attribute (RADIUS attribute 60) in Access-Request packets.

The optional behavior requires that the value of the challenge must be 16 bytes. If the challenge is not 16 bytes long, **authd** ignores the optional configuration and sends the challenge as the CHAP-Challenge attribute.

To configure the optional behavior, you use the **chap-challenge-in-request-authenticator** statement at the **[edit access profile *profile-name* radius options]** hierarchy level.

[See [Configuring RADIUS Server Options for Subscriber Access](#).]

- **NAS-Port-ID string values and order (MX Series)**—Starting in Junos OS Release 15.1, you can specify additional optional information in the NAS-Port-ID (RADIUS attribute 87), which identifies the physical interface used to authenticate subscribers. In addition, you can override the default order in which the optional values appear in the NAS-Port-ID and specify a customized order for the optional values.

You can now include the following additional information when configuring the **nas-port-id-format** statement at the **[edit access profile *profile-name* radius options]** hierarchy level:

- **interface-text-description**—interface's description string
- **postpend-vlan-tags**—VLAN tags using **<outer>-<inner>**

Use the **order** option at the **[edit access profile profile-name radius options nas-port-id-format]** hierarchy level to specify the non-default order in which the optional information appears in the NAS-Port-ID string.

[See [Configuring a NAS-Port-ID with Additional Options.](#)]

- **Changes to LAC connect speed derivation (MX Series)**—Starting in Junos OS Release 15.1, the following changes are made to the methods that specify a source for the LAC to derive values for the Tx-Connect-Speed and Rx-Connect-Speed that it sends to the LNS in AVP 24 and AVP 38:
  - The **static** method is no longer supported for specifying a source, but it is still configurable for backward compatibility. If the **static** method is configured, the LAC falls back to the port speed of the subscriber access interface.
  - The default method has changed from **static** to **actual**.
  - The **actual** method now has the highest preference when multiple methods are configured; in earlier releases, the **ancp** method has the highest preference.
  - When the **pppoe** method is configured and a value is unavailable in the PPPoE IA tags for the Tx speed, Rx speed, or both, the LAC falls back to the port speed. In earlier releases, it falls back to the **static** method.
- **Change to show services l2tp tunnel command (MX Series)**—Starting in Junos OS Release 15.1, the **show services l2tp tunnel** command displays tunnels that have no active sessions. In earlier releases, the command does not display tunnels without any active sessions.
- **Support for LAC sending AVP 46 (MX Series)**—Starting in Junos OS Release 15.1, when the LAC terminates a PPP session, it generates a PPP disconnect cause and includes this information in the PPP Disconnect Cause Code (AVP 46) when it sends a Call-Disconnect-Notify (CDN) message to the LNS. The code value is 0, which indicates a global error with no information available.
- **New option to limit the maximum number of logical interfaces (MX Series routers with MS-DPCs)**—Starting in Junos OS Release 15.1, you can include the **limited-ifl-scaling** option with the **network-services enhanced-ip** statement at the **[edit chassis]** hierarchy level to impose a limitation on the maximum number of logical interfaces on MX Series routers with MS-DPCs to be 64,000 for enhanced IP network services mode. Using the **limited-ifl-scaling** option prevents the problem of a collision of logical interface indices that can occur in a scenario in which you enable enhanced IP services mode and an MS-DPC is also present in the same chassis. A cold reboot of the router must be performed after you set the **limited-ifl-scaling** option with the **network-services enhanced-ip** statement. When you enter the **limited-ifl-scaling** option, none of the MPCs are moved to the offline state. All the optimization and scaling capabilities supported with enhanced IP mode apply to the **limited-ifl-scaling** option.
- **Local DNS configurations available when authentication order is set to none (MX Series)**—Starting in Junos OS Release 15.1R2, subscribers get the DNS server addresses when both of the following are true:
  - The authentication order is set to **none** at the **[edit access profile profile-name authentication-order]** hierarchy level.



- A DNS server address is configured locally in the access profile with the **domain-name-server**, **domain-name-server-inet**, or **domain-name-server-inet6** statement at the **[edit access profile *profile-name*]** hierarchy level.

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

- **Change in support for L2TP statistics-related commands (MX Series)**—Starting in Junos OS Release 15.1R2, statistics-related **show services l2tp** commands cannot be issued in parallel with **clear services l2tp** commands from separate terminals. In earlier releases, you can issue these **show** and **clear** commands in parallel. Now when any of these **clear** commands is running, you must press Ctrl+c to make the **clear** command run in the background before issuing any of these **show** commands. The relevant commands are listed in the following table:

<b>clear services l2tp destination</b>	<b>show services l2tp destination extensive</b>
<b>clear services l2tp session</b>	<b>show services l2tp destination statistics</b>
<b>clear services l2tp tunnel</b>	<b>show services l2tp session extensive</b>
	<b>show services l2tp session statistics</b>
	<b>show services l2tp summary statistics</b>
	<b>show services l2tp tunnel extensive</b>
	<b>show services l2tp tunnel statistics</b>



**NOTE:** You cannot run multiple **clear services l2tp** commands from separate terminals. This behavior is unchanged.

## System Logging

- **System log message for key encryption key (KEK) creation or activation**—Starting with Junos OS Release 15.1, messages similar to the following system log message are generated by the gkmd process when a KEK is created or deleted:

```
root@host> show log messages | grep "Created KEK"
May 16 13:42:01 host gkmd[25450]: Created KEK with SPI {283f0f68 95739eb6 -
37a72054 d775ccde} for group vpn vpn-group6-srx
clear group security on the server:
root@host> show log messages | grep "Deleted KEK"
May 16 14:00:41 host gkmd[25450]: Deleted KEK with SPI {283f0f68 95739eb6 -
37a72054 d775ccde} for group vpn vpn-group6-srx
```

## System Management

---

- **Change to process health monitor process (MX Series)**—Starting in Junos OS Release 15.1R2, the process health monitor process (pmond) is enabled by default on the Routing Engines of MX Series routers, even if no service interfaces are configured. To disable the pmond process, include the **disable** statement at the **[edit system processes process-monitor]** hierarchy level.

## User Interface and Configuration

---

- **Space character not a valid name or value in CLI**—Starting in Junos OS Release 15.1, you cannot create a name or value in the CLI using only single or multiple space characters. Existing configurations that include names or values consisting of only the space character cannot upgrade to Junos OS Release 15.1. The space character can still be used as part of a name or value in the CLI, as long as other characters are present.
- **New flag to control errors when executing multiple RPCs through a REST interface (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, you can stop on an error when executing multiple RPCs through a REST interface by specifying the **stop-on-error** flag in the HTTP POST method.

[See [Submitting a POST Request to the REST API](#).]

- **Changed available REST interface cipher suites when Junos OS is in FIPS mode (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, when Junos OS is in FIPS mode, you can only configure cipher suites with a FIPS-compliant hash algorithm for the REST interface to the device. To configure a cipher suite, specify the **cipher-list** statement at the **[edit system services rest https]** hierarchy level.

[See [cipher-list \(REST API\)](#).]

- **New command to view disk space usage in configuration database (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, you can use the **show system configuration database usage** command to see how much of the disk space is allocated for storing previous versions of the committed configurations and how much space is used by the configuration data.

[See [show system configuration database usage](#).]

## VPNs

---

- **Group VPNv2 member devices allow multiple Group VPNv2 groups to share the same gateway (MX Series)**—In order to make configuration and debugging easier, starting in Junos OS Release 15.1, multiple Group VPNv2 groups can use the same gateway. The commit check for a unique tuple of **<local\_address, remote\_address, routing\_instance>** across groups has been removed. The same tuple is now checked for uniqueness across all gateways. This allows multiple groups to share the same gateway for their Group VPNv2 traffic.

### Related Documentation

- [New and Changed Features on page 21](#)
- [Known Behavior on page 67](#)

- [Known Issues on page 68](#)
- [Resolved Issues on page 76](#)
- [Documentation Updates on page 99](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)
- [Product Compatibility on page 113](#)

## Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R2 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.

- [Hardware on page 67](#)
- [MPLS on page 68](#)
- [Subscriber Management and Services \(MX Series\) on page 68](#)
- [System Logging on page 68](#)

---

### Hardware

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

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

---

## MPLS

- **Removal of SRLG details from the SRLG table only on the next reoptimization of the LSP**—If an SRLG is associated with a link used by an ingress LSP in the router, then on deleting the SRLG configuration from that router, the SRLG gets removed from the SRLG table only on the next reoptimization of the LSP. Until then, the output displays Unknown-XXX instead of the SRLG name and a nonzero srlg-cost of that SRLG for **run show mpls srlg** command.

---

## Subscriber Management and Services (MX Series)

- The **show ppp interface *interface-name* extensive** and **show interfaces pp0** commands display different values for the LCP state of a tunneled subscriber on the LAC. The **show ppp interface *interface-name* extensive** command displays STOPPED whereas the **show interfaces pp0** command displays OPENED (which reflects the LCP state before tunneling). As a workaround, use the **show ppp interface *interface-name* extensive** command to determine the correct LCP state for the subscriber.

---

## System Logging

- **Text string deprecated in syslog messages that are converted to SNMP traps (M Series, MX Series, and T Series)**—In the syslog messages that are converted to SNMP traps for event policies, the "trap sent successfully" text string is deprecated.

### Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 52](#)
- [Known Issues on page 68](#)
- [Resolved Issues on page 76](#)
- [Documentation Updates on page 99](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)
- [Product Compatibility on page 113](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R2 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.

- [Forwarding and Sampling on page 69](#)
- [General Routing on page 69](#)
- [Infrastructure on page 71](#)

- [Interfaces and Chassis on page 71](#)
- [Layer 2 Features on page 72](#)
- [MPLS on page 72](#)
- [Network Management and Monitoring on page 72](#)
- [Platform and Infrastructure on page 73](#)
- [Routing Protocols on page 74](#)
- [Services Applications on page 74](#)
- [Software Installation and Upgrade on page 75](#)
- [Subscriber Access Management on page 75](#)
- [User Interface and Configuration on page 75](#)
- [VPNs on page 75](#)

---

### Forwarding and Sampling

- This defect is seen only when an 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](#)
- 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 subscriber management environment, after scaling subscribers login/logout multiple times, the MX Series routers may hang subscriber in the terminated state and be stuck in the backup accounting queue. The reason is that, when authentication daemon (authd) is trying to fetch data from session database (SDB), error (for example, session not found, or an SDB deadlock or during the SDB recovery period) may occur, and this error may cause the router to fail to notify the client daemon to clean up the service records. In this case, the subscribers may not be able to send Acct-Stop messages to RADIUS server and end up with staying in terminated state. [PR1041070](#)
- When "satop-options" is configured on an E1 with Structure-Agnostic TDM over Packet (SAToP) encapsulation, after Automatic Protection Switching (APS) switchover, some SAToP E1s on the previously protect interface (now working) start showing drops. [PR1066100](#)

- When VMX is deployed, initially there is no management port configuration, so configuration needs to be applied by serial console. The console for VMX is set to 9600 baud rate. With this rate, only a small number of configuration lines can be pasted at a time. [PR1068152](#)
- In subscriber management environment, the PPP daemon (jpppd) might crash repeatedly due to a memory double-free issue. [PR1079511](#)
- In a two members MX Series Virtual Chassis (MXVC) environment, when "set virtual-chassis no-split-detection" is configured, if split master condition happens, which is caused by split events (i.e. loss of all adjacencies by link failure, FPC restarts, chassis power-down, Routing Engine reboots, etc), then once the VCP adjacency is formed again, the current design could not determine best chassis to win the protocol mastership election properly. Instead, only the final election step (that is, choose the member device with the lowest MAC address) is used to elect the master device (protocol master of the VC, or VC-M). [PR1090388](#)
- Starting with Junos OS Release 14.1, Entropy Label Capability is enabled by-default on all Juniper Networks [MX] systems. On PTX transit LSRs that carry LSPs with Entropy Label Capability, packet loss can be observed due to data errors when one or more labeled route entries are not properly removed from the hash table (ie. following LSP optimization or MBB event) because the 'stale' entries are pointing to corrupted route memory. As a result, when the MPLS label that's associated with the 'stale' entry is re-used, data errors are seen for packets using the corresponding label. [PR1100637](#)
- After Junos OS Release 13.3R1, IPCMON infrastructure is added to debug IPCs between PFEMAN and the Routing Engine. When convergence occurs, string processing of IPCMOM will take added time. Then the slow convergence will be seen. It is a performance issue. [PR1100851](#)
- PDB errors thrown when ffp is executed in PDB-unsupported platforms. To avoid these errors, the scope of syslog errors are limited to the PDB-supported platforms based on the error code returned. (PDB\_ERROR\_PLATFORM\_NOT\_SUPPORTED in this case). [PR1103035](#)
- If fpc offline configuration statement is configured after the presence of Non-recoverable faults, then offline action will not be performed. [PR1103185](#)
- cpcdd core observed in scaled scenario. [PR1103675](#)
- After Line Card reboot subscriber traffic is counted against underlying interface instead one created for subscriber. [PR1110493](#)
- Resolved problem with Syslog messages generated like "krt\_decode\_resolve for 239.255.255.250, 101.11.67.33: no logical interface for index 1073741825" when Multicast packets are received on Subscriber interfaces. [PR1110967](#)
- The MS-MPC service card will fail to restart automatically when the POWER\_ZONE it is powered under loses power when running the MX960 with high capacity power supplies split into two separate power zones. [PR1112716](#)
- Under certain conditions, when the Junos OS Routing Engine tries to send an IP packet over an IPIP tunnel, the lookup might end up in an infinite loop between two IPIP tunnels.

This is caused by a routing loop causing the tunnel destination for Tunnel#A to be learned through Tunnel#B and the other way round. [PR1112724](#)

- Under certain conditions, when the Junos OS Routing Engine tries to send an IP packet over a GRE tunnel, the lookup might end up in an infinite loop between two GRE tunnels. This is caused by a routing loop causing the tunnel destination for Tunnel#A to be learned through Tunnel#B and the other way round. [PR1113754](#)

## Infrastructure

- When "show version detail" CLI command has been executed, it will call a separate gstatd process with parameter "-vvX". Because the gstatd could not recognize these parameters, it will run once without any parameter then exit. In result of "show version detail", following information could be seen: user@mx960> show version detail  
 Hostname: mx960 Model: mx960 Junos: 13.3R6-S3 JUNOS Base OS boot [13.3R6-S3]  
 JUNOS Base OS Software Suite [13.3R6-S3] JUNOS Kernel Software Suite [13.3R6-S3]  
 JUNOS Crypto Software Suite [13.3R6-S3] <snipped> file: illegal option -- v usage:  
 gstatd [-N] gstatd: illegal option -- v usage: gstatd [-N] <snipped> At the same time, log lines like following might be recorded in syslog: Aug 25 17:43:35 mx960 file: gstatd is starting. Aug 25 17:43:35 mx960 file: re-initialising gstatd Aug 25 17:43:35 mx960 mgd[14304]: UI\_CHILD\_START: Starting child '/usr/sbin/gstatd' Aug 25 17:43:35 mx960 gstatd: gstatd is starting. Aug 25 17:43:35 mx960 gstatd: re-initialising gstatd Aug 25 17:43:35 mx960 gstatd: Monitoring ad2 Aug 25 17:43:35 mx960 gstatd: switchover enabled Aug 25 17:43:35 mx960 gstatd: read threshold = 1000.00 Aug 25 17:43:35 mx960 gstatd: write threshold = 1000.00 Aug 25 17:43:35 mx960 gstatd: sampling interval = 1 Aug 25 17:43:35 mx960 gstatd: averaged over = 30 Aug 25 17:43:35 mx960 mgd[14304]: UI\_CHILD\_STATUS: Cleanup child '/usr/sbin/gstatd', PID 14363, status 0x4000 Aug 25 17:43:35 mx960 mgd[14304]: UI\_CHILD\_EXITED: Child exited: PID 14363, status 64, command '/usr/sbin/gstatd' [PR1078702](#)
- SSD failure does not trigger Routing Engine switchover. [PR1102978](#)

## Interfaces and Chassis

- The following log can be seen on OTN capable pics after each commit, which indicates incorrect stats TLV setting. No service impact found. /kernel: ge-1/1/0: Unknown TLV type 356 /kernel: ge-1/1/0: Unknown TLV type 361 /kernel: ge-1/1/0: get tlv ppfeid 0xe-0/2/0: get tlv ppfeid 0xe-0/3/0: get tlv ppfeid 0xe-1/2/0: get tlv ppfeid 0xe-1/3/0: get tlv ppfeid 0xe-2/0/0: get tlv ppfeid 0xe-2/1/0: get tlv ppfeid 0xe-2/2/0: get tlv ppfeid 0xe-2/3/0: get tlv ppfeid 0xe-5/1/0: get tlv ppfeid 0xe-5/1/1: get tlv ppfeid 0xe-5/1/2: get tlv ppfeid 0xe-5/1/3: get tlv ppfeid 0xe-5/1/4: get tlv ppfeid 0xe-5/1/5: get tlv ppfeid 0xe-5/1/6: get tlv ppfeid 0xe-5/1/7: get tlv ppfeid 0xe-5/1/8: get tlv ppfeid 0xe-5/1/9: get tlv ppfeid 0 [PR1057594](#)
- The 'optics' option will now display data for VCP ports: show interfaces diagnostics optics vcp-0/0/0 [PR1106105](#)
- On T Series Multichassis platform, when offline and then online the LCC from SCC (e.g. executing the CLI command "set chassis lcc 0 offline" command, and then executing "delete set chassis lcc 0 offline") in quick successions (that is, within the timeout setting for peer to reconnect, 60 seconds, which is not configurable), kernel replication

error "ENOENT" may occur, which can cause ksyncd to crash and in thus trigger a live vmcore. Additionally, this is a timing issue and LCC offline followed by online within 60 seconds is the only known trigger so far. As a workaround, on the safer side, it is recommended to online the LCC after 120 seconds. [PR1108048](#)

---

## 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](#)

---

## MPLS

- RSVP graceful restart does not function for LSPs that have a forwarding adjacency (FA) label-switched path (LSP) as a next hop. [PR60256](#)

---

## 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 TRIO-only filter - count correctly as configured in the filter Now, despite the SNMP mib walk for firewall OIDs lists all OIDs and appropriate values, messages logs will report the following logs for every interface that has this TRIO-only filter applied. > 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) The above 2 issues are addressed in this PR fix. [PR988566](#)
- If Routing Engine protocol mastership is not established and a daemon like mib2d tries to register with shm-rtssdb for ifState updates, it may not receive updates. Due to this, a recent fix was introduced to delay the above registration until Routing Engine's Protocol Mastership is resolved. As a side effect of this fix - we see this core. In this case SNMP Requests have landed on the mib2d, before it has connected to shm-rtssdb and initialized its interface database. As a fix - (To Avoid SNMP Requests landing on mib2d before Routing Engine-Mastership is resolved) we have delayed the MIB registration as well. Hence after Routing Engine bootup, once Protocol Mastership is resolved - mib2d will connect to shm-rtssdb and then register its MIBs with snmpd. So no snmp requests will be received in mib2d until mastership is resolved. [PR1114001](#)



---

## Platform and Infrastructure

---

- 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](#)
- If with both MPC/MSDPC and other type of DPCs equipped, for local switching at mesh group level, split horizon on PW interfaces will not work and this would cause packets to loop back to same PW interface. [PR1084130](#)
- Aggregate Ethernet (AE) interfaces in combination with shared-bandwidth-policer might lead to Packet Forwarding Engine policer corruption if there are child member links configured on the same Packet Forwarding Engine and the AE interface is being reconfigured (add/delete of logical units). This corruption could alter the policer rate programmed in hardware and lead to unexpected policer behavior. A different trigger with physical interface flap illustrating the same symptoms are tracked in PR1035845. [PR1084912](#)
- On MX Series platform, if ingress "multicast-replication" is configured, the throughput of the multicast may get reduced due to unnecessary threads during Packet Forwarding Engine operation. In addition, only the performance of multicast traffic may get influenced (some of the multicast packets may get dropped on the Packet Forwarding Engine) by the issue. This PR has fixed/enhanced the performance. Now the performance limit should only be capped by fabric bandwidth in ingress Packet Forwarding Engine. In addition, before this fix, there was a limitation that VPLS/Bridging can't run with ingress-replication feature as its BUM traffic can't be handled by ingress-replication feature. This PR removed that limitation as well. Now BUM traffic for VPLS/Bridging is following normal multicast replication path even with ingress-replication feature. [PR1089489](#)
- Once the culprit-flows are detected, it might keep reporting culprit-flows logs even after corresponding flows are gone on MPC3E or MPC4E. [PR1102997](#)
- The following fields have been added to v10 Sampling (IPFIX) template and data packets: - SAMPLING RATE - SAMPLING INACTIVE TIMEOUT - SAMPLING ACTIVE TIMEOUT - TOTAL PACKETS EXPORTED - TOTAL FLOWS EXPORTED. [PR1103251](#)
- On MX Series-based platform, in MX Series Virtual Chassis (MXVC) environment, if the subscriber logical interface (IFL) index 65793 is created (for example, when carrying 15K DHCPv4 subscribers to exceed IFL index creation 65793) and the IEEE 802.1p rewrite rule is configured (for example, using CoS rewrite rules for host outbound traffic), due to usage of incorrect logical interface index, the Virtual Chassis Control Protocol Daemon (vccpd) packets (for example, Hello packets) transmission may get lost on all VC interfaces, which may lead to VC decouple (split brain state, where the cluster breaks into separate parts). As a workaround, either delete the rewrite rule (delete class-of-service host-outbound-traffic ieee-802.1 rewrite-rules), or find the logical interface in jnh packet trace that is not completing the vccpd send to other chassis and at Routing Engine clearing that subscriber interface may resolve the issue. [PR1105929](#)

- When "shared-bandwidth-policer" is configured with aggregate Ethernet (AE) has more than one member link on the same Packet Forwarding Engine and the policer is configured with "physical-interface-policer" configuration statement, if reconfiguration occurs (for example, adding/deleting new logical units, logical interface flap...), Packet Forwarding Engine may problem wrong policer during this reconfiguration process, which could ultimately lead to unexpected packet drop/loss within the referenced wrong shared policer. [PR1106654](#)
- DHCP End options (option 255) is missing by DHCP-relay agent (where 20 bytes DHCP options 82 inserted) for client DHCP discover message with 19bytes padding. [PR1110939](#)
- Inline 6rd and 6to4 support for XL and XL-XM based platforms. [PR1116924](#)

---

### Routing Protocols

- 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 this soft-assert-coredump, without any impact to traffic/service. [PR815146](#)
- In rare cases, rpd may write a core file with signature "rt\_notbest\_sanity: Path selection failure on ..." The core is 'soft', which means there should be no impact to traffic or routing protocols. [PR946415](#)
- With this change the default label hold timer was increased for 10 seconds to 60 seconds. [PR1093638](#)

---

### Services Applications

- In the NAT environment, the jnxNatSrcPoolName OID is not implemented in jnxSrcNatStatsTable. [PR1039112](#)
- When polling to jnxNatSrcNumPortInuse via SNMP MIB get, it might not be displayed correctly. [PR1100696](#)
- 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. [PR1109447](#)

## Software Installation and Upgrade

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

## Subscriber Access Management

- When the MX Series 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](#)

## 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

- Refer to release note of [PR535844](#) It is planned for future releases of Junos OS to 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 21](#)
- [Changes in Behavior and Syntax on page 52](#)

- [Known Behavior on page 67](#)
- [Resolved Issues on page 76](#)
- [Documentation Updates on page 99](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)
- [Product Compatibility on page 113](#)

## 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.

- [Class of Service \(CoS\) on page 76](#)
- [Forwarding and Sampling on page 77](#)
- [General Routing on page 80](#)
- [High Availability \(HA\) and Resiliency on page 84](#)
- [Interfaces and Chassis on page 84](#)
- [Layer 2 Features on page 88](#)
- [MPLS on page 89](#)
- [Network Management and Monitoring on page 89](#)
- [Platform and Infrastructure on page 90](#)
- [Routing Policy and Firewall Filters on page 94](#)
- [Routing Protocols on page 94](#)
- [Services Applications on page 96](#)
- [Software Installation and Upgrade on page 97](#)
- [Subscriber Access Management on page 97](#)
- [User Interface and Configuration on page 98](#)
- [VPNs on page 98](#)

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

---

- For an ATM interface configured with hierarchical scheduling, when a traffic-control-profile attached at ifd (physical interface) level and another output traffic-control-profile at ifl (logical interface) level, flapping the interface might crash the FPC. [PR1000952](#)
- 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 times out 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, when aggregate Ethernet (AE) interface is in link aggregation group (LAG) Enhanced mode, after deactivating and then activating one child link of the LAG, the feature that runs on AE interface rather than on the child link (for example, IEEE-802.1ad rewrite rule) may fail to be executed. [PR1080448](#)
- After restarting chassisd or doing an in-service software upgrade from 13.2R8.2 to 13.3R7.3, results in the following messages seen in syslog:  
cosd\_remove\_ae\_ifl\_from\_snmp\_db ae40.0 error 2 Messages appear to be harmless with no functionality impact. [PR1093090](#)
- On MX104 platform, when we configure rate-limit for the logical tunnel (lt-) interface, the commit will fail. As a workaround, we can use firewall filter with policer to achieve the same function. [PR1097078](#)
- 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 (logical interface) adding message and the logical interface 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](#)
- When performing the Routing Engine switchover without GRES enabled, due to the fact that the Class-of-Service process (cosd) may fail to delete the traffic control profile state attached to logical interface (IFL) index, the traffic-control-profile may not get programmed after the logical interface index is reused by another interface. [PR1099618](#)

### Forwarding and Sampling

- When there are no services configured, datapath-traced daemon is not running. In the PIC, the plugin continues to try for the connection and continuous connection failure logs are seen. [PR1003714](#)
- In IP security (IPsec) VPN environment, after performing the Routing Engine switchover, the traffic may fail to be forwarded due to the SAs may not be downloaded to the PIC, or due to some security associations (SAs) on the PIC may incorrectly hold references for old Security Policy Database (SPD) handles while SPD has deleted its entries in the Security Association Database (SAD). [PR1047827](#)
- On all Junos OS based platforms, there are two different types of memory blocks that might be leaked. The first issue is rpd-trace memory block leak. There is one block each for any trace files opened for rpd. They could be leaked for each time a configuration commit is done. Around 40 bytes are leaked per operation. The issue does not occur in Junos OS Release prior to 14.1. The second issue is rt\_parse\_memory block leak which could happen during the configuration of aggregate routes, configuration information might not be freed. Around 16384 bytes are leaked per operation. This issue is a day-1 issue. [PR1052614](#)

- When enabling pseudowire subscribers the "show subscribers extensive" command does not display CoS policies applied to the subscriber interface. This issue was fixed in 13.3R6, 14.1R5 and 14.2R3. [PR1060036](#)
- For MX Series Virtual Chassis (MX-VC) with scaled subscribers, for example, 100K DHCP/20K PPPoE subscribers. If the Virtual Chassis port (VCP) FPCs also house the uplink ports and the "indirect-next-hop-change-acknowledgements" and "krt-nexthop-ack-timeout" configuration statements are configured along with the protection mechanism, after the master Routing Engine in the Virtual Chassis master router (VC-Mm) is powered down, the traffic loss and subscriber loss might be observed due to the indirect next-hop change acknowledgement timeout. With this fix, the upper limit for "krt-nexthop-ack-timeout" is changed from 100 seconds to 250 seconds. [PR1062662](#)
- For MX-VC platform, performing unified ISSU in scaled subscribers environment might cause all VC members to get restarted unexpectedly. [PR1070542](#)
- After rebooting the BNG with scaled subscribers, a dynamic-profile add request might fail, causing bbe-smgd (subscriber management daemon) to crash, then some subscribers might fail to login. [PR1071850](#)
- Juniper Networks device is not sending an error code to the Open vSwitch Database (OVSDb) client when the commit fails. Now a graceful mechanism is introduced to handle netconf configuration errors. If a netconf commit fails, the transaction will be routed to a failed queue. The transaction remains in the failed queue, until the user takes action to explicitly clear the transaction from the failed queue using the CLI. New CLI commands to show and clear failed netconf transactions. `user@router> show ovsdb netconf transactions Txn ID Logical-switch Port VLAN ID 1 vlan100 user@router> clear ovsdb netconf transactions` [PR1072730](#)
- On MX Series-based platform, when the Layer 3 packets destined to an Integrated Routing and Bridging (IRB) interface and then hit the underlying Layer 2 logical interfaces (IFLs), due to the egress feature list of the Layer 2 logical interfaces may get skipped, the features under the family bridge (for example, the firewall filter) on the Layer 2 interfaces may not be executed. [PR1073365](#)
- The issue is seen while moving an interface from one mesh group to another. [PR1077432](#)
- In scaled subscriber management environment (for example, 3.2K PPPoE subscribers), after heavy login/logout, the session setup rate keeps decreasing and also PAP-NAK messages are sent with "unknown terminate code". This continues till Broadband Network Gateway (BNG) does not accept PPP sessions and all newly incoming sessions are stuck in PAP Authentication phase (No PAP ACK received). [PR1075338](#)
- The license-check process may consume more CPU utilization. This is due to a few features trying to register with the license-check daemon which license-check would not be able to handle properly and results in high CPU on Routing Engine. Optimization is done through this fix, to handle the situation gracefully so that high CPU will not occur. [PR1077976](#)
- From Junos 14.1R1, if the hidden configuration statement "layer-4 validity-check" is configured, the Layer4 hashing will be disabled for fragmented IP traffic. Due to a

defect, the Multicast MAC rewrite is skipped in this case, the fragmented multicast packets will be sent with incorrect destination MAC. [PR1079219](#)

- 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](#)
- OTN based SNMP Traps such as jnxFruNotifOperStatus and jnxIfOtnNotificationOperStatus are raised by offline/online MIC although no OTN interface is provisioned. [PR1084602](#)
- Invalid Ethernet Synchronization (ESMC) frames may be transmitted by MX router when activating LAG and tag-protocol-id under interfaces. [PR1084606](#)
- On a device with lt and ams interfaces configured, walking ifOutOctets or other similar OID's may cause a "if\_pfe\_ams\_ifdstat" message to print. This is a cosmetic debug-level entry, which was incorrectly set to critical-level. [PR1085926](#)
- In the specific configuration of a LT interface in a VPLS instance and the peer-unit of this LT interface configured with family inet6 using vrrp, the kernel may crash when the FPC is online. [PR1087379](#)
- 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](#)
- In rare cases, SSH or telnet traffic might hit incorrect filter related to SCU (Source Class Usage) due to the defect in kernel filter match. This issue comes when the filter has match condition on source class ID. [PR1089382](#)
- In rare cases, MX Series routers might crash while committing inline sampling related configuration for INET6 Family only. [PR1091435](#)
- In a fib-localization scenario, IPv4 addresses configured on service PICs (SP) will not appear on FIB-remote FPCs although all local (/32) addresses should, regardless of FIB localization role, install on all Packet Forwarding Engines. There is no workaround for this and it implies that traffic destined to this address will need to transit through FIB-local FPC. [PR1092627](#)
- There are entries for PEM in jnxFruEntry in VMX. It is not necessary and is cosmetic. [PR1094888](#)
- When route convergence occurred, the new gateway address is not updated correctly in inline-jflow route-record table (route-record table is used by sampling), and the sampling traffic forwarding might be affected, but normal routing would be not affected. [PR1097408](#)
- After upgrading to Junos OS Release 14.1R1 and higher, loopback ISO family address may be stuck in KRT queue. [PR1097778](#)
- When BGP multipath is enabled in a Virtual Routing and Forwarding (VRF), if "auto-export" and "rib-group" are configured to leak BGP routes from this VRF table to another, for example, the default routing table, then traffic coming from the default routing instance might not be properly load balanced due to the multipath-route leaked

into the default routing table is not the active route. This is a random issue. As a workaround, only use "auto-export" to exchange the routes among the routing tables. [PR1099496](#)

## General Routing

---

- There is hardware design flaw with 2x10GE MIC and 4x10GE MIC today which introduces +/-6.2ppm frequency offset for SyncE operation. In order to correct this, the framing of the PIC and interface has to be matched (which will not be by default). [PR932659](#)
- SNMP MIB walk of object "jnxSpSvcSet" gives hardcoded value as "EXT-PKG" for SvcType. [PR1017017](#)
- With Multiservices MPCs (MS-MPCs) or Multiservices MICs (MS-MICs) installed on MX Series platform, when trying to view the Network Address Translation (NAT) mappings for address pooling paired (APP) and/or Endpoint Independent Mapping (EIM) from a particular private or a public IP address, all the mappings will be displayed. [PR1019739](#)
- On MX Series router with MPC3E/MPC4E/MPC5E/MPC6E if the Packet Forwarding Engine has inline NAT configured or is processing inline GRE decapsulation with packet-sizes between 100B-150B, in some very corner cases, traffic blackhole might be seen due to incorrect cell packing handling. On T4000 with FPC type 5, when these cards are processing any packets sizes between 133B-148B in certain sequences causes incorrect cell packing handling. [PR1042742](#)
- Queue stats on LSQ interfaces are not properly cleaned up when queuing enabled on the IFD and the queues hosted at IFD level. This happens when a subsequent delete and create of LSQ interface (not always though) - 14.1R4.10. [PR1044340](#)
- MPC with Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC (MIC-3D-4COC3-1COC12-CE) might crash. This problem is very difficult to replicate and a preventive fix will be implemented to avoid the crash. [PR1050007](#).
- In subscriber management environment, the Berkeley Database (DB) may get into deadlock state. It is brought on by multiple processes attempting to simultaneously access or update the same subscriber or service record. In this case, due to the access to DB were 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 request 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](#)
- With inline L2TP IP reassembly feature configured, the MX Series routers with MPCs/MICs might crash due to a memory allocation issue. [PR1061929](#)
- In subscriber management environment, if IPv6 family is not enabled in the dynamic profile, the IPv6 Router Advertisement message will not be sent through the dynamic subscriber interface. As a workaround, you can enable family inet6 in the dynamic profile. [PR1065662](#)
- When setting the syslog to debug level (any any), you may note reoccurring messages of the form "ifa for this rt ia is not present, consider ifa as ready". These messages are



logged for IPv6 enabled interfaces when receiving forwarded packets and cause no harm. Set a higher debug level to avoid seeing them. [PR1067484](#)

- The static route prefers the directly connected subnet route for resolving the nexthop rather than performing a longest prefix match with any other available routes. In case of longest prefix route being desired in customer deployment, it will result in traffic loss issue. Now a new configuration statement "longest-match" is introduced to enable longest prefix matching behavior when desired: set routing-options static route <destination prefix> next-hop <address> resolve longest-match. [PR1068112](#)
- In subscriber management environment, changing the system time to the past (for example, over one day) may cause the processes (for example, pppoe, and autoconfd) that use the time to become unresponsive. [PR1070939](#)
- 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](#)
- 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 MX-based MPC, low priority threads which monitor various things in the background on a periodic basis such as voltage, temperature, stats counters, hardware status exist. 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 it is a non impacting issue. [PR1071408](#)
- Traffic throughput test between MPC1/1E/2/2E card and MPC2E/3E NG card, the flowing from MPC1/1E/2/2E card to MPC2E/3E NG card is lesser then from MPC2E/3E NG card to MPC1/1E/2/2E card. [PR1076009](#)
- Vendor provided the fix, which includes conditional check. [PR1076369](#)
- In a Q-in-Q setup, if outer vlan tag is coming with EtherType 0x88a8, it is not possible to create dynamic vlan interface on Junos 13.1X42 or 14.1X51 releases. [PR1080734](#)
- On MX Series platform with MS-MPC/MS-MIC, in some mspmand process crash scenarios, after the mspmand coredump is finished or almost finished, PIC kernel also crashes and dumps vmcore. The mspmand cores in these scenario are readable but vmcores are not. [PR1081265](#)
- In DHCPv6 prefix delegation over PPPoE scenario, when forwarding the control packet from the Routing Engine to the DHCPv6 identity association for prefix delegation (IA\_PD) address over PPPoE, for instance, executing ping from Routing Engine targeting the client's PD address, the traffic may get dropped on the device. [PR1081579](#).
- If a router has Service PIC equipped but without any Service PIC specific configurations, the CPU usage on this PIC/FPC might be high. Have some configurations under below configuration statement could prevent from this issue: [system processes process-monitor traceoptions] OR [chassis fpc <fpc slot> pic <pic slot> adaptive-services service-package extension-provider] OR [services] [PR1081736](#)

- In multi-homing and signal active EVPN scenario, if IRB interface is included in the instance, when the DF-CE link flaps, due to a timing issue, the DF might send L3 EVPN routes with label 0 to remote PEs, causing traffic to be dropped at remote PE. [PR1082287](#)
- 'show interfaces queue <ifl>' stats are not correct with RLSQ warm-standby mode. Issue seen on MPCs and MICs as well in 14.1R4.10. [PR1082417](#)
- TCP messages do not have their MSS adjusted by the Multiservices MIC and MPC if they do not belong to an established session. [PR1084653](#)
- With a scaled subscribers system, repeatedly doing tcpdump of subscriber interface and press ctrl+c might cause bbe-smgd daemon memory growing, which will in turn causing crash, SDB corruption and some other daemons crashing. Following signs may be seen when this problem is hit: log messages like: "/kernel: cmd bbe-smgd pid 1997 tried to use non-present sched\_yield" tcpdump stops working bbe-smgd no longer accepts new vty sessions. [PR1085944](#)
- In some rare conditions, depending on the order in which configuration steps were performed or the order in which hardware modules were inserted or activated, if PTP master and PTP slave are configured on different MPCs on MX Series router acting as BC, it might happen that clock is not properly propagated between MPCs. This PR fixes this issue. [PR1085994](#)
- 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](#)
- mspmand.core is observed while making ms-mic offline with IPsec and Jflow configured on same ms-mic with dynamic IPSEC tunnels. [PR1086819](#)
- If the ALG is receiving UDP fragmented control traffic (e.g. SIP control packets) continuously, the mspmand process (which manages the service PIC) might crash due to buffer error. [PR1087012](#)
- Due to a change in an existing PR, group names in the configuration must be a string of alphanumericals, dashes or underscores. There is not workaround other than following the group name instructions. [PR1087051](#)
- On LAC (L2TP Access Concentrator) router with session client-idle-timeout configured, the tunneled PPP session will always keep active due to the PPP control messages are accounting as user data. [PR1088062](#)
- Wrong ESH checksum computation with non-zero Ethernet Padding in Juniper MX Series router. [PR1091396](#)
- The mspmand process might crash due to prolonged flow-control with TCP ALGs under the following possible scenario, mostly when the following conditions happen together: 1. When the system is overloaded with TCP ALG Traffic 2. There are lots of retransmissions and reordered packets. [PR1092655](#)
- When the control path is busy/stuck for service PIC, the AMS member interface hoisted by it might be down, but when the busy/stuck condition is cleared, the member interface might not recover, and AMS bundle still shows the PIC as inactive. [PR1093460](#)

- On TCP ALG, if there are a lot of retransmissions and reordered TCP packets, and the system is overloaded due to the TCP traffic, the mspmand (which manages the service PIC) process might crash. [PR1093788](#)
- In a scaled Broadband Subscriber Management environment (in this case, 16K subscribers), when Access Node Control Protocol (ANCP) CoS adjustment is configured, the minimum rate instead of the shaping-rate might be wrongly applied to some subscribers and causes traffic loss. [PR1094494](#)
- Extensive Header integrity checks will be done for packets which match a service set which has NAT/SFW configured. 1. Enable Header integrity checks by default when SFW or NAT is configured in same service set. This is inline with ukernel behavior 2. Retain the configuration statement for use by other plugins such as IPsec which may want to enforce header integrity if needed 3. Ensure that the cmd "show services service-sets statistics integrity-drops" works if sfw/nat is configured [PR1095290](#)
- The issue is because of the software problem. Just after the system reboots, rpd process is determining the Routing Engine mastership mode too early before chassisd is determining the mastership, which would cause overload feature to not work properly. [PR1096073](#)
- If a service-PIC is configured to simultaneously function as both an MS interface and as a member of an AMS interface, then some settings under services-options may not apply correctly. These settings are A) syslog\_rate\_limit, B) fragment-limit, C) reassembly-timeout and D) jflow\_log\_rate\_limit. [PR1096368](#)
- For Junos 13.3R1 and later, the DPC card might experience a performance degradation when it's transferring bidirectional short packets (64B) in inline rate. [PR1098357](#)
- Some of the new revisions (for example, REV 30, REV 31) of the MICs can not come up with NG-MPC2 or NG-MPC3 line card. We can check the MIC version by CLI command "show chassis hardware detail | no-more". root@user> show chassis hardware detail | no-more Hardware inventory: Item Version Part number Serial number Description .. FPC 2 REV 14 750-054901 CADJ3871 MPC3E NG PQ & Flex Q CPU REV 11 711-045719 CADN5465 RMPC PMB MIC 0 REV 30 750-028392 CAEB9203 3D 20x 1GE(LAN) SFP <<<<<REV>[PR1100073](#)
- When the null pointer of jbuf is accessed (jbuf, that is, a message buffer is allocated only when the packet is ready to process. The buffer is freed after the packet completes ALG handling is accessed), for example, when using the Microsoft Remote Procedure Call (MS RPC) (as observed, issue may also happen on Sun Microsystems RPC) Application-level gateway (ALG) with NAT (stateful firewall is used as a part of the service chain), if the traffic matching configured universal unique identifier (UUID) is arrived on the ALG, the mspmand (which manages the Multiservice PIC) crash occurs. [PR1100821](#)
- In broadband edge (BBE) environments, for example, if the interface-set is created corresponding to SVLAN, then multiple logouts and logins will create a new interface-set index. When the interface-set index range goes above 65535, executing CLI command "show interfaces interface-set queue egress" will cause 100% CPU usage. As a workaround, we can use the specified interface-set name instead of using the wildcard. [PR1101648](#)

- On MX dual Routing Engine platforms, if there are a large number of addresses (in this case, there are > 500 addresses configured, the issue might be observed around 472 addresses) configured on lo0.0, when the Broadband Edge subscriber management daemon (bbe-smgd) replicating these addresses to the standby Routing Engine, the internal 8K replication buffer may get exceeded. Due to this failure, memory leak (around 45MB every time error is encountered) may occur when bbe-smgd tries to delete the object. Since lo0.0 object gets created/destroyed over and over, bbe-smgd runs out of memory and crash eventually. [PR1101652](#)
- 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](#)
- On MX Series platform, the output of CLI command "show system subscriber-management route" may be shown as empty. [PR1104808](#)
- An IPv4 filter configured to use the filter block with term that has both "from precedence" and another non 5-tuple (i.e. not port, protocol, address) will cause an XL/EA based board to reboot. Example: set firewall family inet filter FILTER fast-filter-lookup set firewall family inet filter FILTER term TERM from precedence PRECEDENCE set firewall family inet filter FILTER term TERM from tcp-established. [PR1112047](#)

---

### High Availability (HA) and Resiliency

- On dual Routing Engine platforms with NSR enabled, when committing scaling configuration (for example, deactivating 500 logical interfaces and performing commit, then activating 500 logical interfaces and commit, the process may need to be performed 3-6 times) to the device, the master Routing Engine would be busy processing commit, due to which the backup does not get data or keepalive from master. In this situation, the protocols (for example, OSPF, or LDP) may get down on the backup Routing Engine due to keepalive timeout. [PR1078255](#)

---

### Interfaces and Chassis

- Chap Local-name default to 8 characters. Should be 32. [PR996760](#).
- If a subscribers-facing AE interface has link protection enabled, offline the primary child link hosted FPC might cause some subscribers to down. [PR1050565](#)
- dcd will crash if targeted-distribution applied to ge ifd via dynamic-profile. [PR1054145](#)
- 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 setting set together on pm3393 chip driver software, so it used to so happen that even before the copy of stats to shadow registers happened, clear was happening which used to go unaccounted. [PR1056232](#)
- When a dynamic PPPoE subscriber with targeted-distribution configured on a dynamic vlan demux interface over aggregated ethernet, the device control daemon (dcd) process might crash during a commit if the vlan demux has mistakenly been removed. The end users cannot visit internet after the crash. This is a rare issue and not easy to be reproduced. [PR1056675](#)

- It is observed that the syslog messages related to kernel and Packet Forwarding Engine may get generated at an excessive rate, especially in subscriber management environment. Most of these messages may appear repeatedly, for example, more than 1.5 million messages may get recorded in 2 hours, and there are only 140 unique messages. Besides, these messages are worthless during normal operation and due to the excessive rate of log generation, it results in high Routing Engine CPU consumption (for example, Routing Engine CPU utilization can be stuck at 100% for a long time (minutes or hours), it depends on the activity of subscribers (frequency of logins and logouts) and on the AI scripts used by the customer) by event process (eventd) might be observed on the device. [PR1056680](#)
- In subscriber management environment, PPP client process (jpppd) might crash as a result of a memory allocation problem. [PR1056893](#).
- 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](#)
- For transit traffic on INLINE LSQ redundancy (rlsq) interface, the input firewall-filter counters are logging zero packet count regardless of traffic flow. Output filter counters are logging correctly. For host-bound traffic, the firewall output counter will get double accounted on Classical rlsq and triple accounted on INLINE rlsq. This issue is targeted to be fixed in Junos 14.1R5. [PR1060659](#)
- In scaling PPP subscriber environment, when the device is under a high load condition (for example, high CPU utilization with 90% and above), the long delay in session timeout may occur. In this situation, the device may fail to terminate the subscriber session (PPP or PPPoE) immediately after three Link Control Protocol (LCP) keepalive packets are missed. As a result, the subscriber fails in reconnect due to old PPP session and corresponding Access-Internal route are still active for some time. In addition to this, it is observed that the server is still sending KA packets after the session has timed out. [PR1060704](#)
- For Junos OS Release 13.3R1 or above, after multiple (e.g. 26) iterations of graceful Routing Engine switchover (GRES), the TNP address of management interface might be deleted incorrectly during switchover, this leads to all FPCs being offline. [PR1060764](#)
- Link Up/Down SNMP traps for AE member links might not be generated, but the SNMP traps for the AE bundle works well. [PR1067011](#)
- In PPP subscriber management environment, the jpppd process might crash for a timing issue. [PR1074545](#)
- When the Ethernet Link Fault Management (LFM) action profile is configured, if there are some errors (refer to the configuration, for example, frame errors or symbol errors) happening in the past (even a long past), due to the improper handling of error stats fetching from kernel, the LFM process (lfmd) may generate false event PDUs and send false alarm to the peer device. [PR1077778](#)
- On MX Series Virtual Chassis (MX-VC) platform, due to a timing issue, the physical interface (ifd) on the same Modular Interface Card (MIC) with Virtual Chassis port

(VCP) might not be created or takes a very long time to be created after rebooting the hosted Modular Port Concentrator (MPC). [PR1080032](#)

- MAX-ACCESS value has been changed in jnx-otn.mib for the following oids:  
jnxOtnIntervalOdu15minIntervalNumber jnxOtnIntervalOtu15minIntervalNumber  
jnxOtnIntervalOtuFec15minIntervalNumber The value has been changed from read-only to not-accessible to be inline with newer MIBs. [PR1080802](#)
- On MX Series platform acting as broadband network gateway (BNG), in Point-to-Point Protocol (PPP) scenario, when using the Internet Protocol version 6 Control Protocol (IPv6CP) for negotiation, if the router receives an IPv6CP Configure-Request packet from client, MX BNG sends the Configure-Request packet, but does not send IPv6CP Configure-Ack packet, in case it does not receive the Configure-Ack that responding to the Configure-Request packet it sent. The behavior does not follow the RFC 1661, which demands both the actions Send-Configure-Request (i.e. IPv6CP-ConfReq from MX to client) and Send-Configure-Ack (i.e. IPv6CP-ConfAck from MX to client) to be conducted on the router without any significant delay. [PR1081636](#)
- With Non-MX Series/service DPCs which are not supported with enhanced-ip, when these unsupported DPCs are in the chassis, the user switches to enhanced-ip and reboots the router, the router should come back up and the unsupported DPCs should stay powered off and not log any alarms. In this case, the non-supported DPCs stay powered off, but they are also continuing to raise alarms. There are two workarounds for this issue; first, power down the FPC prior to changing enhanced-ip mode; second, perform a hard restart by "restart chassis-control immediately" to restore. Both of these workarounds will impact traffic through the router. [PR1082851](#)
- In MX virtual chassis (MXVC) scenario, during unified ISSU operation, the new master Routing Engine does not have the MXVC SCC's system MAC address. It just has its local system MAC address. The address is not replicated between local Routing Engines, and the new master Routing Engine is not yet connected to the MXVC SCC to receive it. Hence, the possibility of overwriting the FPC with an address that does not match the previous address exists. [PR1084561](#)
- The VRRP preempt hold time is not being honored during NTP time sync and system time is changed. [PR1086230](#)
- On MX Series Virtual Chassis (MX-VC) platform with "subscriber-management" enabled, after power up/reboot, the VC backup router (VC-B) experiences a rapid sequence of role transitions from no-role to VC master router (VC-M) to VC-B, the expected local GRES and a reboot of the former master Routing Engine might not happen on the VC-B. Some of the FPCs on it might be stuck in "present" state and eventually rebooted. [PR1086316](#)
- Deactivating/activating logical interfaces may cause BGP session flapping when BGP is using VRRP VIP as the 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](#)
- 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](#)

- When an interface on SFPP module in MIC is set disabled, after pulling out the SFPP and then insert it, the remote direct connected interface might get up unexpectedly. [PR1090285](#).
- After removing a child link from AE bundle, in the output of "show interface <AE> detail", the packets count on the remaining child link spikes, then if add back the previous child link, the count recover to normal. [PR1091425](#)
- For Junos OS version 14.1X51-D60 or 14.1X50-D105, when DHCP local server is configured, the DHCP subscribers might be unable to come up. [PR1092553](#)
- In MX Series Virtual Chassis (MXVC) environment, when rebooting the system or the line cards which contain all the Virtual Chassis port (VCP) links, because line cards might fail to complete the rebooting process within 5 minutes, the timer (that is, the amount of time allowed for the LCC to connect to the SCC) started by the master router might expire which may cause the VCP links establishment failure. In addition, this issue is not specific to the line cards type, based on the observation, the timer (5 minutes) may expire on a MX2020 with all 20 FPCs equipped as well. [PR1095563](#)
- On PB-20C12-ATM2-SMIR PIC, port 0 and port 1 are configured with clock source as external, if Loss of signal (LOS) is inserted on port 0, the port 0 will go down, the expected behavior is clock being used from port 1. But in this case, port 0 down will results in port 1 flapping and reporting SONET phase lock loop (PLL) errors. [PR1098540](#)
- In VRRP environment, with VRRP configured over double tagged interface and VRRP delegate-processing enabled, the PDUs are generated with only one tag and the outer tag is not added, because of which, the PDUs will get dropped at the receiving end. The similar configuration that may cause the issue might be seen as below, .. protocols { vrrp { delegate-processing; <<<<< "delegate-processing" is enabled for VRRP } ... interfaces { xe-0/0/3 { flexible-vlan-tagging; unit 0 { vlan-tags outer 2000 inner 200; <<<<< VRRP is configured over double tagged interface family inet { address 10.10.10.147/29 { vrrp-group 17 { virtual-address 10.10.10.145; priority 100; accept-data; } } } } } .. [PR1100383](#)
- After configuring related ae interface configuration, we might find some of ae interfaces disappear in MX-VC. It seemed that ae interfaces are not allocated MAC address from chassisd properly. \* This issue only happens in the first configuration timing after rebooting/restarting chassisd. So even if you configure related ae interface configuration repeatedly, you cannot find this issue. When this issue happens these message will be seen in the messages logs. -----  
lab@router\_re0> show log messages| match CHASSISD\_MAC\_ADDRESS\_AE\_ERROR  
Jun 26 16:04:34.064 router\_re0 scchassisd[2008]:  
CHASSISD\_MAC\_ADDRESS\_AE\_ERROR: chassisd MAC address allocation error for ae4 Jun 26 16:04:34.105 router\_re0 /kernel: Jun 26 16:04:34.064 router\_re0 scchassisd[2008]: CHASSISD\_MAC\_ADDRESS\_AE\_ERROR: chassisd MAC address allocation error for ae4 ----- Restore ae interfaces \* This is not workaround. deactivate/activate ae interfaces. (We need to do this to all disappeared ae interfaces.) [PR1100731](#)



- 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](#)
- Due to the fact that the error injection rate configured by user on Routing Engine via CLI command "bert-error-rate" may not be programmed in the hardware register, the PE-4CHOC3-CE-SFP, PB-4CHOC3-CE-SFP, MIC-3D-4COC3-1COC12-CE, and MIC-4COC3-1COC12-CE-H may fail to inject bit errors during a Bit Error Ratio Test (BERT). [PR1102630](#)
- On MX240 or MX480 platform with at least two DC modules (PN: 740-027736) equipped, when shutting down one of the PEMs and then turn it on again, even the PEM is functioning, the "PEM Fan Fail" alarm might be observed on the device due to software logic bug. There is no way to clear the ALARM\_REASON\_PS\_FAN\_FAIL for I2C\_ID\_ENH\_CALYPSO\_DC\_PEM once it has been raised. [PR1106998](#)

---

## Layer 2 Features

- Under rare circumstances it is possible for the DHCP drop counts for reason SEND ERROR to be incremented twice for a single failure. [PR1009296](#)
- MTU change is not advised on the Ethernet ring protection (ERP) ring interfaces unless ring is in idle condition. Changing ring interface MTU while ring is not in idle state might result in change in the forwarding state of the interface which can lead to loop in the ring. [PR1083889](#)
- When family bridge was configured and committed, l2ald repeated restarting with core. After l2ald repeated restarting several times, it stopped working due to thrashing condition. Core of l2ald will be seen with the following configuration. set interfaces fxp0 unit 0 family bridge interface-mode access set interfaces fxp0 unit 0 family bridge vlan-id 100 When the configuration is committed, message like following is logged and core is generated. l2ald[1624]:  
../../../../src/junos/usr.sbin/l2ald/l2ald\_vpls\_flood.c:3117: insist '!err' failed l2ald[1734]:  
../../../../src/junos/usr.sbin/l2ald/l2ald\_vpls\_flood.c:3117: insist '!err' failed l2ald[1769]:  
../../../../src/junos/usr.sbin/l2ald/l2ald\_vpls\_flood.c:3117: insist '!err' failed l2ald[1993]:  
../../../../src/junos/usr.sbin/l2ald/l2ald\_vpls\_flood.c:3117: insist '!err' failed l2ald[2195]:  
../../../../src/junos/usr.sbin/l2ald/l2ald\_vpls\_flood.c:3117: insist '!err' failed ... init:  
l2-learning is thrashing, not restarted [PR1089358](#)
- During interface flaps, a high amount of TCN (Topology Change Notification) might get propagated causing other switches to get behind due to high amount of TCN flooding. This problem is visible after the change done from 11.4R8 onwards which propagates TCN BPDU immediately and not in the pace of the 2 second BPDU. Hello interval to speed up topology change propagation. The root cause is that the TCNWHILE timer of 4 seconds is always reset upon receiving TCN notifications causing the high churn TCN propagation. [PR1089580](#)



- In MX Series Virtual Chassis (MXVC) environment, when packets come from a interface (for example, xe-16/0/1.542) situated on one member of VC (for example, VC member 1), if the ingress Packet Forwarding Engine (for example, FPC16 PFE0, who runs hash to determine which interface it should send the packet to) decides that it should send the packet via another interface (for example, xe-4/0/1.670) situated on different member (for example, VC member 0), it will send the frame to member 0 via the vcp-intf. In case of xe-4/0/1.670 belongs to an AE bundle which has multiple child links, a hash need to be run on Packet Forwarding Engine carrying the VCP port (receiving side on member 0) to determine which one is the egress Packet Forwarding Engine within member 0 to send the packet out after vcp-intf gets the packet. This hash result should get the same result as the ingress Packet Forwarding Engine. If it is not the case, then the packet would get dropped on Packet Forwarding Engine on member 0. [PR1097973](#)
- With scaled subscribers connected, restarting one of MPCs might cause subscribers unable to log in for about 2 minutes. [PR1099237](#)

## MPLS

- 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, and the RSVP local repair might fail to process more than 200 LSPs at one time, the traffic might 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](#)
- In race conditions, the rpd process on backup Routing Engine might crash when BGP routes are exported into LDP by egress-policy and configuration changes during the rpd process synchronizing the state to backup rpd process. [PR1077804](#)
- On dual Routing Engine platform with GRES, the kernel synchronization process (ksyncd) might crash on the backup Routing Engine when adding of route pointing to indirect nexthop on system. [PR1102724](#)

## Network Management and Monitoring

- In rare cases, when the mib2d process attempts connection with the snmpd process and there are pending requests waiting to be finished, the mib2d process might crash and the CPU utilization is high around the same time as the crash happens. [PR1076643](#)
- Due to a bug in jnxIfcInline mib, a high order interface churn such as the one done by the submitter in this case, can lead to a mib2d core. The situation is recovered after the core and no other impact is seen. [PR1105438](#)

## Platform and Infrastructure

---

- Adaptive load-balance functionality is only supported for unicast traffic. If the aggregate bundle contains logical interfaces for bridge or vpls domains, flooded traffic might get dropped. [PR821237](#)
- In dual Routing Engines scenario with NSR configuration, the configuration statement "groups re0 interfaces fxp0 unit 0" is configured. If disable interface fxp0, backup Routing Engine is unable to proceed with commit processing due to SIGHUP not received, the rpd process on backup Routing Engine might crash. [PR974430](#)
- When Network Configuration Protocol (NETCONF) service is used on the device, after the NETCONF session is established, because all the output that contain <error> tag might be incorrectly converted into <rpc error>, the management daemon (mgd) may crash on the device. As the following example, the output that contains <error> tag may lead to the crash. user@re0> show subscribers address 1000 | display xml .. <error junos:style="input-error"> <<<<<< The output contain <error> tag and may trigger the crash. [PR975284](#)
- On MX Series Virtual Chassis (MX-VC) platform, mirroring of OAM packets may not work as expected if the OAM packet is traversing through multiple Packet Forwarding Engines (for example, the mirrored port and VCP port are on separate Packet Forwarding Engines). [PR1012542](#)
- In EVPN scenario, MPC may crash with core-file when any interface is deleted and add that interface to an aggregated Ethernet bundle or changing the ESI mode from all-active to single-active. [PR1018957](#)
- LSI logical interface input packet and byte stats are also added to core logical interface stats, but when the LSI logical interface goes down and the core logical interface stats are polled, there is a dip in stats. The fix is to restore LSI logical interface stats to core logical interface before deleting the LSI logical interface. [PR1020175](#)
- Under very rare situations, Packet Forwarding Engines on the following linecards, as well as the compact MX80/40/10/5 series, may stop forwarding transit traffic: - 16x10GE MPC - MPC1, MPC2. This occurs due to a software defect that slowly leaks the resources necessary for packet forwarding. Interfaces handled by the Packet Forwarding Engine under duress may exhibit incrementing 'Resource errors' in consecutive output of 'show interfaces extensive' output. A Packet Forwarding Engine reboot via the associated linecard or chassis reload is required to correct the condition. [PR1058197](#)
- On MX Series router with frame-relay (FR) CCC to connect FR passport devices. If some of the FR circuits carry traffic without any valid FR encapsulations, the MX Series based Packet Forwarding Engine drops those frames. [PR1059992](#)
- If a Radius server is configured as accounting server, when it is non-reachable, the auditd process might be stressed with huge number of audit logs to be sent to the accounting server, which might cause auditd to crash. [PR1062016](#)
- Modifying IEEE-802.1ad rewrite-rule on the fly might be unable to change IEEE-802.1p ToS values for inner VLAN in QinQ. [PR1062817](#)

- In Junos release 13.3R6 or 14.2R3, for PPPoE subscribers over the aggregated Ethernet (ae) interface, the output of "show interface statistics <pp> detail" command shows the ingress/egress traffic statistics for the aggregate interface instead of the statistics for PP/DEMUX logical interface. [PR1069242](#)
- Having "shared-bandwidth-policer" on an aggregated ethernet interface; if a member interface flapped, the NPC which the interface belongs may restart. Similar issue may also happen when changing the firewall policer configuration. [PR1069763](#)
- When Integrated routing and bridging (IRB) interface is configured with Virtual Router Redundancy Protocol (VRRP) in Layer 2 VPLS/bridge-domain, in corner cases after interface flapping, MAC filter ff:ff:ff:ff:ff:ff is cleared from the Packet Forwarding Engine hardware MAC table, so the IRB interface may drop all packets with destinations MAC address FFFF:FFFF:FFFF (e.g. ARP packet). [PR1073536](#)
- It tries to check allotted power for all the FPCs, here in the CHASSISD\_I2CS\_READBACK\_ERROR logs it shows for the FPCs which are not present in chassis. It just calls i2cs\_readback() to read i2c device and fails there as these FPCs? slots are blank and prints those readback errors. Also the errors are harmless: "CHASSISD\_I2CS\_READBACK\_ERROR: Readback error from I2C slave for FPC" Fix: Code to check 'if power has been allotted to this FPC', needs to be executed only if the FPC is present. [PR1075643](#)
- When using the "ping detail" command, the interface number is provided on the output instead of the interface name. [PR1078300](#)
- During a unified in-service software upgrade (ISSU), DHCP control traffic (renew/rebinds) might be dropped on ingress Packet Forwarding Engine. [PR1079812](#)
- When an MX chassis network-services is "enhanced-ip" and an AE is part of a Layer 2 bridge (bridge-domain or VPLS), there is a possibility that an incorrect forwarding path might be installed causing traffic loss. This could happen when first applying the configuration, restarting the system or restarting the line card. [PR1081999](#)
- On MX Series-based platform, the "RPF-loose-mode-discard" feature is not working when configured within a Virtual Router routing instance. The feature is working only when configured in the main instance. [PR1084715](#)
- With MSDPC equipped on BNG, there might be a memory leak in ukernel, which eventually causes MSDPC to crash and restart. [PR1085023](#)
- In Junos OS Releases 13.3R3, 14.1R1, 14.2R1, there is a new feature, an extra TLV term is added to accommodate the default action for the "next-interface" when the corresponding next-interface is down. While doing a unified ISSU from an image without the feature to an image with this feature, all MPCs might crash. [PR1085357](#)
- If there are scaling unicast routes (e.g. 500k) in NG-MVPN VRF, and the provider-tunnel is PIM, when PIM on PE has multiple upstream neighbors and any of them could be its rpf neighbor, performing GRES/NSR Routing Engine switchover might cause multicast traffic loss due to the different view of rpf neighbor between the master Routing Engine and the slave Routing Engine. [PR1087795](#)
- The prompt for SSH password changed in Junos OS Release 13.3, from "user@host's password:" to "Password:". This change breaks the logic in "JUNOS/Access/ssh.pm"

which is located in /usr/local/share/perl/5.18.2/ on Ubuntu Linux, for example.

[PR1088033](#)

- On MX Series router with MPC1/1E, MPC2/2E line cards in a broadband edge environment with scaled (in this case 250K) subscribers, the FPC heap (dynamic memory) utilization increases significantly during an in-service software upgrade (ISSU). [PR1088427](#)
- On MX Series platform with MPC/MIC or T4000 FPC5, TCP session with MS-Interface/AMS-Interface, configuration is not established successfully with the "no-destination-port" or "no-source-port" configuration statements configured under forwarding-options hierarchy level. [PR1088501](#)
- Issue is specific to 64-Bit RPD and config-groups wildcard configuration specific 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? (i.e. 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 below 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](#)
- On MX Series router, if ifl (logical interface) is configured with VID of 0 and parent ifd (physical interface) with native-vlan-id of 0, when sending L2 traffic received on the ifl to Routing Engine, the VID 0 will not be imposed, causing the frames to get dropped at Routing Engine. [PR1090718](#)
- When an interface on MQ-based FPC is going to link down state, in-flight packet on interface transmit path will be stuck on the interface and never drained until the interface comes up again. As a result, small number of such stacked packets will be sent out when the interface is going to UP state. No other major impact should be seen after those packets are drained. [PR1093569](#)
- On MX2020/2010 router, an SPMB core file will be seen if there are bad XF chips (fabric chip) on SFB, which might trigger Routing Engine/CB switchover. [PR1096455](#)
- In 64-bit Junos OS environment, the Representational State Transfer (REST) API fails to start when configured with "set system services rest ...". [PR1097266](#)
- When a P2MP LSP is added or deleted at ingress LSR, traffic loss is seen to existing sub-LSP(s) at transit LSR which replicates and forwards packet to egress PEs. This issue only affects MX Series based line card. [PR1097806](#)
- The "shared-bandwidth-policer" configuration statement is used to enable configuration of interface-specific policers applied on an aggregated Ethernet bundle to match the effective bandwidth and burst-size to user-configured values. But this feature is broken from Junos OS Release 14.1R1 when "enhanced-ip" is configured on MX Series platform with pure MX Series-based line cards. The bandwidth/burst-size

of policers attached to Aggregated Ethernet interfaces are not dynamically updated upon member link adding or deletion. [PR1098486](#)

- On MX Series-based platform, when the type of the IPv6 traffic is non-TCP or non-UDP (for example, next header field is GRE or No Next Header for IPv6), if the traffic rate is high (for instance, higher than 3.5Mpps), the packet re-ordering may occur. [PR1098776](#)
- On MX Series-based line cards, when the prefix-length is modified from higher value to lower value for an existing prefix-action, heap gets corrupted. Due to this corruption, the FPC might crash anytime when further configurations are added/deleted. The following operations might be considered as a workaround: Step 1. Delete the existing prefix-action and commit Step 2. Then re-create the prefix-action with newer prefix-length. [PR1098870](#)
- In an MPLS L3VPN network with a dual-homed CE router connected to different PE routers, a protection path should be configured between the CE router and an alternate PE router to protect the best path. When BFD is enabled on the BGP session between the CE and the primary PE router, with local traffic flowing from another CE connected with the primary PE to this CE, after bringing the interface down on the best path, the local repair will be triggered by BFD session down, but it might fail due to a timing issue. This will cause slow converge and unexpected traffic drop. [PR1098961](#)
- When the BFD is running on multi LU (lookup chip) Packet Forwarding Engine (such as MPC3 or MPC4), incoming BFD packet might be processed with a firewall filter on different logical-routers's loopback interface. If the firewall filter is discarding/rejecting BFD, the packets will be dropped incorrectly. [PR1099608](#)
- On MX Series-based platform, before creating a new unicast nexthop, there is a check to see if there is at least 512k DoubleWords (DW) free. So, even the attempting NH requires only a small amount of memory (for example, < 100 DWs), if there is no such enough free DWs (that is, 512k), the check will fail and the end result is that the control plane will quit adding this NH prematurely - stopping at ~80% of capacity. With the fix, it will check for 64k free DWs which is lower reference watermark for available resource, thereby ensuring that can allocate resource. [PR1099753](#)
- From Junos OS Release 14.1 and above, IPv6 mobility packets with Heartbeat option that the length of the mobility header (including the ethernet encapsulation and main IPv6 header) extends beyond 128 Bytes will be discarded as bad IPv6 option packet due to a logic error in packet handling. [PR1100442](#)
- On MX Series platform, when using the 64-bit image, if the configuration statement "source-address" is configured for the "radius-server" as the following, the RADIUS request may not be sent to RADIUS server due to the failure of setting the "source-address" on the device. `user@re0> show configuration system radius-server .. source-address 10.1.1.1; <<<<<` The configuration statement that may cause the issue [PR1103517](#)

- On T4000 platform with FPC Type-5 equipped, after performing unified ISSU, due to the fact that only 6 out of 16 temperature sensors may get initialized, the temperature reading for the line card may be shown as "Absent". [PR1104240](#)
- Due to a software defect found in 13.3R7.3 and 14.1R5.4, Juniper Networks strongly discourage the use of Junos OS software version 13.3R7.3 on routers with MQ-based MPC. This includes MX Series with MPC1, MPC2; all mid-range MX Series. [PR1108826](#)

---

### Routing Policy and Firewall Filters

- In Class-of-Service (CoS) environment, there is a possibility (happened twice so far and not reproducible in the lab) that routing protocol process (rpd) may crash because the CoS memory may get incorrectly freed and then allocated again. [PR1062616](#)
- On the platform that M7i/M10i with enhanced CFEB, M320 with E3-FPC, M120, and MX with DPC, when the flood filter is configured in VPLS instance on the Packet Forwarding Engine, if the Packet Forwarding Engine receives a filter change (for example, FPC reboot occur and comes up), the line card may fail to program the filter. [PR1099257](#)

---

### Routing Protocols

- 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 PIM multicast-only fast reroute (MoFRR) environment, when issuing CLI command "show multicast route extensive" on egress edge router, due to missing null check while showing label information for reverse-path forwarding (RPF) nexthop, an error might be seen in the output of the command. In addition, the routing protocol process (rpd) may crash on the device. [PR983140](#).
- 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 mutli-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](#).
- There are two issues in the PR: (1) In multicast environment, Incoming interface list (IIF) list has only RPF interface, designated forwarder (DF) winners are not added in the list in backup Routing Engine. (2) "Number of downstream interfaces" in show pim join extensive is not accounting Pseudo-VXLAN interface. [PR1082362](#)
- If a policy statement referred to a routing-table, but the corresponding routing instance is not fully configured (ie. no instance-type), commit such configuration might cause the rpd process to crash. [PR1083257](#).
- 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](#)

- When there are a number of secondary BGP routes in inet.0, an SNMP walk of inet.0 by the bgp4 MIB can cause a core if the corresponding primary routes are being deleted. [PR1083988](#).
- 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](#)
- When BGP route is leaked to a routing-instance and there is an import policy to overwrite the route preference, if damping is also configured in BGP, the BGP routes which were copied to second table cannot be deleted after routes were deleted in master table. This is a day-1 issue. [PR1090760](#)
- When removing BGP Prefix-Independent Convergence (PIC) from the configuration, the expected behavior is that any protected path would become unprotected. But in this case, the multipath entry that contains the protection path (which is supposed to be removed) remains active, until BGP session flaps or the route itself flaps. As a workaround, we can use "commit full" command to correct or to commit. [PR1092049](#)
- In BGP environment, when configuring RIB copy of routes from primary routing table to secondary routing table (for example, by using the CLI command "import-rib [ inet.0 XX.inet.0]") and if the second route-table's instance is type "forwarding", due to the BGP routes in secondary routing table may get deleted and not correctly re-created, the routes may be gone on every commit (even commit of unrelated changes). As a workaround, for re-creating the BGP routes in secondary route table, use CLI command "commit full" to make configuration changes. [PR1093317](#)
- In Junos OS Release 9.1 and later, RFC 4893 introduces two new optional transitive BGP attributes, AS4\_PATH and AS4\_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. In this case, when AS4\_AGGREGATOR attribute (18) is received from a 2-byte AS peer (note AS4\_AGGREGATOR attribute is only received when the aggregator has 4-byte AS but this peer only supports 2-byte AS), NSR synchronization with standby Routing Engine would fail, causing session constantly bouncing on standby Routing Engine (hogging CPU). [PR1093615](#)
- The rpd process might crash when resolve-vpn and rib inet.3 are configured under separate levels (BGP global, group and peer). The fix is if anybody configures a family at a lower level, reset the state created by either of configuration statements from higher levels. This behavior conforms with our current behavior of family configuration - which is that any configuration at a lower level is honored and the higher level configuration is reset. [PR1094499](#).
- When BGP routes has multiple protocol nexthops including discard/reject and other IGP nexthops, the discard/reject nexthop will be selected as BGP nexthop, which will cause traffic loss. [PR1096363](#)



- When polling SNMP OID isisPacketCounterTable 1.3.6.1.2.1.138.1.5.3, the rpd process might crash. [PR1101080](#).
- When the IS-IS configurations have been removed, the IS-IS LSDB contents get flushed. If at the same time of this deletion process, there is an SPF execution (that is, try to access the data structures at same time when/a fraction of seconds after freeing its content), routing protocol process (rpd) crash occurs. [PR1103631](#)

### Services Applications

---

- When an MX Series router configured as an LNS sends an Access-Request message to RADIUS for an LNS subscriber, the LNS now includes the Called-Station-ID-Attribute when it receives AVP 21 in the ICRQ message from the LAC. [PR790035](#)
- In IPsec environment, after performing the Routing Engine switchover (for example, performing Graceful Routing Engine Switchover) or chassis reboot (that is, whole device is powered down and powered UP again), due to the key management daemon (kmd) may be launched before the Routing Engine mastership is finalized, it may stop running on the new master Routing Engine. [PR863413](#)
- In CG-NAT or statefull firewall environment, due to a null pointer check bug, the MS-DPC might crash every few hours. Note that this is a regression issue. [PR1079981](#)
- The crash happens if in a http flow, the flow structure is allocated at a particular memory region. There is no workaround but the chances of hitting this issue are very low [PR1080749](#)
- On Layer 2 Tunnel Protocol (L2TP) network server (LNS), during L2TP session establishment, when receiving Incoming-Call-Connected (ICCN) messages with Last Sent LCP CONFREQ Attribute Value Pair (AVP) but without Initial Received LCP CONFREQ and Last Received LCP CONFREQ AVPs, the jl2tpd process might crash. [PR1082673](#)
- On Trivial File Transfer Protocol (TFTP) Application Layer Gateway (ALG) with NAT translation type "dynamic-nat44" configured, MS-DPC/MS-MPC/MS-MIC might crash when processes the TFTP packets. [PR1091179](#)
- On M Series platform, in Layer 2 Tunneling Protocol (L2TP) network server (LNS) environment, not all attributes (Missing NAS-Identifier, NAS-Port-Type, Service-Type, Framed-Protocol attributes) within Accounting-Request packet are sending to the RADIUS server. [PR1095315](#)
- If MS-DPC is used in CG-NAT environment, in a very rare condition, when the MS-DPC tries to delete a NAT mapping entry (e.g. entry timeout), error might occur and the MS-DPC might get rebooted and then generate a core file. [PR1095396](#)
- Some values of MIB object jnxSrcNatStatsEntry might be doubled when AMS (or rsp) interface and NAT are configured together. [PR1095713](#)



---

## Software Installation and Upgrade

---

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

## Subscriber Access Management

---

- In subscriber management environment, after deactivating a service with Change of Authorization (CoA) dynamic requests, if the Acct-Stop response is not received, the Broadband Network Gateway (BNG) will send CoA NAK message when the same service is activated again. The authd process crash will be observed and some sessions are stuck and cannot be terminated after terminating sessions. [PR1004478](#)
- The authd process memory leaks slowly when subscribers login and logout, which eventually leads the process to crash and generate a core file. [PR1035642](#)
- On MX Series routers, the generic authentication service process (authd) may fail to send Acct-off message to the RADIUS server. This is because management daemon (mgd) might not notify the authd prior to executing system reboot or system shutdown. Also, the authd might fail to generate the Acct-off message as well when it is terminated and there are no active subscribers. [PR1053044](#)
- In subscriber management environment with Remote Authentication Dial In User Service (RADIUS) server configured, when performing scaling subscribers login/logout, the device may stuck in RADIUS communication. [PR1070468](#)
- In subscriber management environment, when dual-stack service is activated by the Change of Authorization (CoA) request from the Radius Server, both families will be activated in the same profile response. Due to a software defect, the service accounting session id is not generated properly and the Service Accounting Messages and Interim-updates failed to be sent out. [PR1071093](#)
- Subscriber is not coming up when CISCO AVPair VSA value is returned in Radius ACCESS-ACCEPT packets in certain scenarios. [PR1074992](#)
- A CoA Request containing LI attributes cannot contain any non-LI service activations, de-activations or variable modifications. [PR1079036](#)
- If authentication-order is configured as none under access profile and domain-name servers (DNS) are configured locally under access profile, then the subscriber will login but will not get DNS addresses which were configured locally. [PR1079691](#)
- In scaled DHCP subscribers environment, the authd process might crash and generate a core file after clearing DHCP binding or logout subscribers. [PR1094674](#)

## 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 not workaround other than following the group name instructions. [PR1087051](#)

## VPNs

---

- Problem, trigger and symptom: On dual Routing Engines, if mvpn protocol itself is not configured, and non stop routing is enabled, the show command "show task replication" on master Routing Engine will list 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](#)
- In PIM Draft-Rosen Multicast VPN (MVPN) environment, in a setup where active C-PR, standby C-RP, C-receivers, C-source are located in different VPN site of MVPN instance, once the link to active C-RP is flapped, PE which connects to C-receivers would send (\*g) join and (s,g,rpt) prune towards standby C-RP, when the PE which connects to standby C-RP receives the (\*g) join and (s,g, rpt) prune over mt-, it ends up updating the (s,g) forwarding entry with mt- as downstream, which is already the incoming interface (IIF). This creates a forwarding loop due to missing check if IIF is same as OIF when PIM make-before-break (MBB) join load-balancing feature is enabled and as a result traffic gets looped back into the network. Loop once formed will remain at least for 210 seconds till the delayed prune timer expires. After this, IIF is updated to the interface towards standby C-RP finally. [PR1085777](#)
- In NG-MVPN spt-only mode with a PE router acts as the rendezvous point (RP), if there are only local receivers, the unnecessary multicast traffic continuously goes to this RP and dropped though it is not in the shortest-path tree (SPT) path from source to receiver. [PR1087948](#)
- When there are more than 2000 outgoing interfaces (OIFs) for a same multicast group on MVPN egress PE, the multicast forwarding entries installed by MVPN might have duplicated OIFs and resulting in duplicated traffic. [PR1095877](#)
- In Internet multicast over an MPLS network by using next-generation Layer 3 VPN multicast (NG-MVPN) environment, when rib-groups are configured to use inet.2 as RPF rib for Global Table Multicast (GTM, internet multicast) instance, the ingress PE may fail to add P-tunnel as downstream even after receiving BGP type-7 routes. In addition, this issue only affects GTM. [PR1104676](#)

### Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 52](#)
- [Known Behavior on page 67](#)
- [Known Issues on page 68](#)
- [Documentation Updates on page 99](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)

- [Product Compatibility on page 113](#)

## Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R2 documentation for the M Series, MX Series, and T Series.

- [Adaptive Services Interfaces Feature Guide for Routing Devices on page 99](#)
- [Broadband Subscriber Sessions Feature Guide on page 100](#)
- [Broadband Subscriber VLANs and Interfaces Feature Guide on page 100](#)
- [High Availability Feature Guide on page 100](#)
- [IPv6 Neighbor Discovery Feature Guide for Routing Devices on page 101](#)
- [Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices on page 101](#)
- [MPLS Applications Feature Guide for Routing Devices on page 102](#)
- [Overview for Routing Devices on page 103](#)
- [Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices on page 103](#)
- [Security Services Administration Guide for Routing Devices on page 103](#)
- [User Access and Authentication Guide for Routing Devices on page 103](#)
- [VPNs Library for Routing Devices on page 103](#)

### Adaptive Services Interfaces Feature Guide for Routing Devices

- In the topic “Inline 6rd and 6to4 Configuration Guidelines,” the next-to-last bullet should state:

Bandwidth for traffic from the 6rd tunnel is limited by the available Packet Forwarding Engine bandwidth; bandwidth for traffic to the 6rd tunnel is limited by the internal VRF loopback bandwidth. SI-IFD loopback bandwidth configuration under the **[edit chassis]** hierarchy has no impact on the 6rd loopback bandwidth.

- The “Configuring Secured Port Block Allocation,” “port,” and “secured-port-block-allocation” topics should include the following note:



**NOTE:** If you make any configuration changes to a NAT pool that has secured port block allocation configured, you must delete the existing NAT address pool, wait at least 5 seconds, and then configure a new NAT address pool. We also strongly recommend that you perform this procedure if you make any changes to the NAT pool configuration, even if you do not have secured port block allocation configured.

## Broadband Subscriber Sessions Feature Guide

---

- In the *Broadband Subscriber Sessions Feature Guide*, the **show network-access aaa radius servers** command topic includes a table that describes the output fields for the command. The table entry for the Status field does not clearly explain when a request starts and ends.

The following information has been added to the NOTE in that table entry: For the purpose of marking a server as **Down** (DEAD), the request includes the original request and any retries that are configured. The 10-second timeout period starts after the initial request and all retries have expired without receiving a response from the server.

The amount of the timeout period that elapses before the server is marked **Down** is not always exactly 10 seconds, and can vary depending on how frequently subscribers are logging in. When subscribers are continually and rapidly logging in, the server is marked as **Down** at 10 seconds. However, if subscribers are logging in less frequently and at a slower pace, then the server is not marked **Down** until a subsequent subscriber attempts to log in. For example, if the subsequent subscriber logs in a minute after the request and all retries lapse, and the 10-second timeout starts, the actual time until the server is marked **Down** is 50 seconds after the timeout starts (the 1 minute between subscriber login minus the 10-second timeout).

## Broadband Subscriber VLANs and Interfaces Feature Guide

---

- The “show subscribers” topic does not fully describe the **vlan-id *vlan-id*** option. This option displays information about active subscribers using a VLAN where the VLAN tag matches the specified VLAN ID. The topic fails to mention that these subscriber VLANs can be either single-tagged or double-tagged. The command output includes information about subscribers using double-tagged VLANs when the inner VLAN tag matches the specified VLAN ID. The command output does not distinguish between these two types of subscribers.

To display only subscribers where the specified value matches only double-tagged VLANs, use the **stacked-vlan-id *stacked-vlan-id*** option to match the outer VLAN tag instead of the **vlan-id *vlan-id*** option.

## High Availability Feature Guide

---

- The following information belongs in the “Nonstop Active Routing Concepts” topic:

If you have NSR configured, it is never valid to issue the **restart routing** command in any form on the NSR master Routing Engine. Doing so results in a loss of protocol adjacencies and neighbors and a drop in traffic.

- The following information belongs in the “Configuring Nonstop Active Routing” topic:

If the routing protocol process (rpd) on the NSR master Routing Engine crashes, the master Routing Engine simply restarts rpd (with no Routing Engine switchover), which impacts routing protocol adjacencies and neighbors and results in traffic loss. To prevent this negative impact on traffic flow, configure the **switchover-on-routing-crash** statement at the **[edit system]** hierarchy level. This configuration forces an NSR Routing Engine switchover if rpd on the master Routing Engine crashes.

- The "Nonstop Active Routing System Requirements" topic should include the **inet-mvpn** and **inet6-mvpn** protocol families for BGP in the list of supported family types. The topic previously documented that NSR supports next-generation MVPN starting with Junos OS 14.1R1, but didn't include the specific names of the next-generation MVPN protocol families in the list.
- The topic "Improving the Convergence Time for VRRP" failed to include the following information:
  - Disable duplication address detection for IPv6 interfaces—Duplicate address detection is a feature of the Neighbor Discovery Protocol for IPv6. Duplicate address detection is enabled by default and determines whether an address is already in use by another node. When duplicate address detection is enabled, convergence time is high after an IPv6 interface that has been configured for VRRP tracking comes up. To disable duplicate address detection, include the **ipv6-duplicate-addr-transmits 0** statement at the **[edit system internet-options]** hierarchy level. To disable duplicate address detection only for a specific interface, include the **dad-disable** statement at the **[edit interfaces interface-name unit logical-unit-number family inet6]** hierarchy level.

### IPv6 Neighbor Discovery Feature Guide for Routing Devices

- The *Secure Neighbor Discovery Guide for Routing Devices* is merged with the *IPv6 Neighbor Discovery Feature Guide for Routing Devices*. We have consolidated these guides and restructured the content in a linear format. The new seamless guide provides related information in a single location for easy navigation and faster access.  
[See [IPv6 Neighbor Discovery Feature Guide for Routing Devices](#).]
- The "NDP Cache Protection Overview," "Configuring NDP Cache Protection," "Example: Configuring NDP Cache Protection to Prevent Denial-of-Service Attacks," and "nd-system-cache-limit" topics failed to include the EX Series, M Series, PTX Series, and T Series as supported platforms. These platforms, as well as the MX series, are all supported.

### Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices

- The Options section for the **flow-export-rate** statement under the hierarchy **[edit forwarding-options sampling instance instance-name family inet output inline-jlow]** did not include the default value. The default value is:  
**Default:** 1 for each Packet Forwarding Engine on the FPC to which the sampling instance is applied.
- The following topics fail to state that for passive monitoring on MX Series routers with MPCs, the **pop-all-labels** statement at the **[edit interfaces interface-name]** hierarchy level pops all labels by default and the **required-depth** statement is ignored.
  - "pop-all-labels"
  - "required-depth"
  - "Enabling Passive Flow Monitoring"

- The "Configuring RPM Timestamping" topic failed to mention that RPM timestamping is also supported on the MS-MPCs and MS-MICs on MX Series routers.
- The description for the **max-packets-per-second**, **maximum-packet-length**, and **run-length** statements at the **[edit forwarding-options sampling instance *instance-name* input]** hierarchy level failed to include the following:



**NOTE:** This statement is not supported when you configure inline flow monitoring (by including the **inline-jflow** statement at the **[edit forwarding-options sampling instance *instance-name* family (inet | inet6) output]** hierarchy level).

- The default value for the **ipv6-flow-table-size** statement at the **[edit chassis fpc *slot-number* inline-services ipv6 flow-table-size]** hierarchy level should state the following:

"If the number of units is not specified, 1024 flow entries are allocated for IPv6."

---

### MPLS Applications Feature Guide for Routing Devices

---

- The "Configuring Miscellaneous LDP Properties," "Configuring the Authentication Key Update Mechanism for BGP and LDP Routing Protocols," "authentication-key-chain (LDP)," and "authentication-key-chain (BGP and BMP)" topics should include the following information: You must also configure the authentication algorithm using the **authentication-algorithm *algorithm*** statement. This statement must be included at the **[edit protocols (bgp | ldp)]** hierarchy level when you configure the **authentication-key-chain *key-chain*** statement at the **[edit protocols (bgp | ldp)]** hierarchy level.
- The "Path Computation for LSPs on an Overloaded Router" topic should state that when you set the overload bit on a router running IS-IS, only new LSPs are prevented from transiting through the router. Any existing Constrained Path Shortest First (CPSF) LSPs remain active and continue to transit through the router. The documentation incorrectly states that any existing LSPs transiting through the router are also rerouted when you configure the overload bit on an IS-IS router.

The topic should also include the following information about bypass LSPs: When you set the overload bit on an IS-IS router, new and existing bypass LSPs are recalculated only when a different event triggers a path recalculation. For example, if you set the smart optimize timer with the **smart-optimize-timer** statement, the bypass LSP is re-routed away from the overloaded router only after the specified time elapses. Otherwise, the bypass LSP continues to transit the overloaded router.

### Overview for Routing Devices

---

- The "Configuring Automatic Mirroring of the CompactFlash Card on the Hard Disk Drive" and the "mirror-flash-on-disk" topics should not include support for MX5, MX10, and MX40 3D Universal Edge Routers. On the MX Series, this feature is supported only on the MX104, MX240, MX480, MX960, MX2010, and MX2020 routers.

### Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices

---

- The table in the "Firewall Filter Nonterminating Actions" topic failed to mention that we recommend that you do not use the nonterminating firewall filter action **next-hop-group** with the **port-mirror-instance** or **port-mirror** action in the same firewall filter.

### Security Services Administration Guide for Routing Devices

---

- The "Distributed Denial-of-Service (DDoS) Protection Overview" topic for Routing Devices has been updated to describe the built-in login overload protection mechanism that is available on MX Series routers.

The login overload protection mechanism (also called a load-throttling mechanism) monitors the incoming subscriber login packets and admits only what the system is capable of handling in accordance with the prevailing load on the system. Packets in excess of what the system can handle are discarded. By shedding this excess load, the system is able to maintain optimal performance and prevent any degradation of login-completion rate under overload conditions. This mechanism uses minimal resources and is enabled by default; no user configuration is required.

The protection provided by this mechanism is secondary to what distributed denial-of-service (DDoS) protection provides as a first level of defense against high rates of incoming packets. DDoS protection operates on the Packet Forwarding Engine and protects against all packet types of all protocols. In contrast, the login overload protection mechanism is located on the Routing Engine and specifically operates only on incoming connection-initiation packets such as DHCPv4 DHCPDISCOVER, DHCPv6 SOLICIT, and PPPoE PADI packets.

### User Access and Authentication Guide for Routing Devices

---

- The "Example: DHCP Complete Configuration" and "dchp" topics should not include support for the MX Series Universal Edge 3D Routers. This feature is supported only on the M Series and the T Series.

### VPNs Library for Routing Devices

---

- The “Routing Instances Overview” topic should include the following instance types: Ethernet VPN (EVPN) and Internet Multicast over MPLS. Use the Ethernet VPN instance type, which is supported on the MX Series only, to connect a group of dispersed customer sites using a Layer 2 virtual bridge. Use the Internet Multicast over MPLS instance type to provide support for ingress replication provider tunnels to carry IP multicast data between routers through an MPLS cloud, using MBGP or next-generation MVPN.

To configure an EVPN instance type, include the **evpn** statement at the **[edit routing-instances *routing-instance-name* instance-type]** hierarchy level. To configure an Internet Multicast over MPLS instance type, include the **mpls-internet-multicast** statement at the **[edit routing-instances *routing-instance-name* instance-type]** hierarchy level.

#### Related Documentation

- [New and Changed Features on page 21](#)
- [Changes in Behavior and Syntax on page 52](#)
- [Known Behavior on page 67](#)
- [Known Issues on page 68](#)
- [Resolved Issues on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 104](#)
- [Product Compatibility on page 113](#)

## 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).

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 15.1 on page 105](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) on page 106](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 107](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 109](#)
- [Upgrading a Router with Redundant Routing Engines on page 109](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 109](#)
- [Upgrading the Software for a Routing Matrix on page 111](#)
- [Upgrading Using Unified ISSU on page 112](#)
- [Downgrading from Release 15.1 on page 112](#)

### Basic Procedure for Upgrading to Release 15.1

When upgrading or downgrading Junos OS, always use the **jinstall** package. Use other packages (such as the **bundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **jinstall** package and details of the installation process, see the [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-x86-32-15.1R2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-15.1R2.9-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 15.1 **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 MX Series routers running Junos OS Release 15.1.

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

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:  
<http://www.juniper.net/support/downloads/>
2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.
4. Select the **Software** tab.
5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new **jinstall** package on the routing platform.



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

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

```
user@host> request system software add validate reboot  
source/jinstall-15.1R2.9-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-15.1R2.9-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 15.1 **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 11.4, 12.3, and 13.3 are EEOL releases. You can upgrade from Junos OS Release 11.4 to Release 12.3 or even from Junos OS Release 11.4 to Release 13.3. 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 12.1 (a non-EEOL release) to Junos OS Release 13.3 or directly downgrade from Junos OS Release 13.3 to Junos OS Release 12.1.

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 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 21](#)
- [Changes in Behavior and Syntax on page 52](#)
- [Known Behavior on page 67](#)
- [Known Issues on page 68](#)
- [Resolved Issues on page 76](#)
- [Documentation Updates on page 99](#)
- [Product Compatibility on page 113](#)

## Product Compatibility

- [Hardware Compatibility on page 113](#)

### 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/>

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

---

These release notes accompany Junos OS Release 15.1R2 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 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Known Issues on page 122](#)
- [Resolved Issues on page 124](#)
- [Documentation Updates on page 127](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

### New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R2 for the PTX Series.

- [High Availability and Resiliency \(HA\) on page 115](#)
- [Interfaces and Chassis on page 115](#)
- [IPv6 on page 116](#)
- [Junos OS XML API and Scripting on page 116](#)
- [Management on page 117](#)
- [MPLS on page 118](#)
- [Routing Protocols on page 118](#)
- [User Interface and Configuration on page 119](#)
- [VPNs on page 120](#)

## High Availability and Resiliency (HA)

- **Unified ISSU support for P2-10G-40G-QSFPP PIC and P2-100GE-OTN PIC (PTX5000)**—Starting with Junos OS Release 15.1, unified in-service software upgrade (ISSU) is supported on P2-10G-40G-QSFPP PIC and on P2-100GE-OTN PIC on FPC2-PTX-P1A FPC in PTX5000 routers. Unified ISSU enables you to upgrade from an earlier Junos OS release to a later one with no disruption on the control plane and with minimal disruption of traffic.

## Interfaces and Chassis

- **Support for including Layer 2 overhead in interface statistics (PTX Series)**—Starting in Junos OS Release 15.1, support is added to account for the Layer 2 overhead size (header and trailer) for both input and output interface statistics in PTX Series routers.
- **Support for dual-rate speed (PTX Series)**—Starting in Junos OS Release 15.1, support for dual rate for the 24-port 10-Gigabit Ethernet PIC (P1-PTX-24-10GE-SFPP) enables you to switch all port speeds to either 1-Gigabit Ethernet or 10-Gigabit Ethernet. The default is 10 Gbps. All ports are configured to the same speed; there is no mixed-rate-mode capability. You can use either the SFP-1GE-SX or the SFP-1GE-LX transceiver for 1 Gbps. Changing the port speed causes the PIC to reboot.

To configure all ports on the P1-PTX-24-10GE-SFPP to operate at 1 Gbps, use the **speed 1G** statement at the `[edit chassis fpc fpc-number pic pic-number]` hierarchy level. To return all ports to the 10-Gbps speed, use the **delete chassis fpc *fpc-number* pic *pic-number* speed 1G** command.

[See [speed \(24-port and 12-port 10 Gigabit Ethernet PIC\)](#) and [10-Gigabit Ethernet PIC with SFP+ \(PTX Series\)](#).]

- **Support for mixed-rate aggregated Ethernet bundles and per-port pseudowire CoS classification on P2-10G-40G-QSFPP PIC and P2-100GE-OTN PIC (PTX5000)**—Starting with Junos OS Release 15.1, you can perform the following actions on the P2-10G-40G-QSFPP PIC and the P2-100GE-OTN PIC on PTX5000 routers:
  - Configure the member links with any combination of rates—10-Gigabit Ethernet, 40-Gigabit Ethernet, and 100-Gigabit Ethernet—on an aggregated Ethernet bundle thereby enabling egress unicast traffic load balancing based on the egress link rate.
  - Classifying port-based pseudowire class of service (CoS) classification which includes Layer 3 IPv4, IPv6, and MPLS classification for interfaces with ethernet-ccc encapsulation.
- **Synchronous Ethernet support for P2-10G-40G-QSFPP PIC, P2-100GE-CFP2 PIC, and P2-100GE-OTN PIC (PTX5000)**—Starting with Junos OS Release 15.1, synchronous Ethernet is supported on the P2-10G-40G-QSFPP PIC, P2-100GE-CFP2 PIC, and P2-100GE-OTN PIC on FPC2-PTX-P1A FPC in PTX5000 routers. Synchronous Ethernet (ITU-T G.8261 and ITU-T G.8264) is a physical layer technology that functions regardless of the network load and supports hop-by-hop frequency transfer, where all interfaces on the trail must support synchronous Ethernet. It enables you to deliver

synchronization services that meet the requirements of the present-day mobile network, as well as future LTE-based infrastructures.

- **CFP-100GBASE-ZR (PTX Series)**—In Junos OS Release 13.3R6, 14.1R4, 14.2R3, and 15.1R1 and later, the CFP-100GBASE-ZR transceiver provides advanced dual polarization-quadrature phase shift keying (DP-QPSK) coherent digital signal processing (DSP) and forward error correction (FEC)-enabled robust tolerance to optical impairments and supports 80 km reach over single-mode fiber. The transceiver is not specified as part of IEEE 802.3 but is built according to Juniper Networks specifications. The following interface module supports the CFP-100GBASE-ZR transceiver:
  - 100-Gigabit Ethernet PIC with CFP (P1-PTX-2-100GE-CFP)

For more information about the interface modules, see the “Cables and Connectors” section in the [PTX Series Interface Module Reference](#).

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#) and [Supported Network Interface Standards by Transceiver for PTX Series Routers](#).]

---

## IPv6

- **Support for outbound-SSH connections with IPv6 addresses (PTX Series)**—Starting with Release 15.1, Junos OS supports outbound-SSH connections with devices having IPv6 addresses.

[See [outbound-ssh](#), [Configuring Outbound SSH Service](#), and [Establishing an SSH Connection for a NETCONF Session](#).]

---

## Junos OS XML API and Scripting

- **Support for replacing patterns in configuration data within NETCONF and Junos XML protocol sessions (PTX Series)**—Starting in Junos OS Release 15.1, you can replace variables and identifiers in the candidate configuration when performing a `<load-configuration>` operation in a Junos XML protocol or NETCONF session. The **replace-pattern** attribute specifies the pattern to replace, the **with** attribute specifies the replacement pattern, and the optional **upto** attribute indicates the number of occurrences to replace. The scope of the replacement is determined by the placement of the attributes in the configuration data. The functionality of the attribute is identical to that of the **replace pattern** configuration mode command in the Junos OS CLI.

[See [Replacing Patterns in Configuration Data Using the NETCONF or Junos XML Protocol](#).]

- **Junos OS SNMP scripts to support custom MIBs (PTX Series)**—Starting with Junos OS Release 15.1, you can use Junos SNMP scripts to support custom MIBs until they are implemented in Junos OS. SNMP scripts are triggered automatically when the SNMP manager requests information from the SNMP agent for an object identifier (OID) that is mapped to an SNMP script for an unsupported OID. The script acts like an SNMP subagent, and the system sends the return value from the script to the network management system (NMS).

[See [SNMP Scripts Overview](#).]

## Management

---

- **Support for enforcing RFC-compliant behavior in NETCONF sessions (PTX Series)**—Starting in Junos OS Release 15.1, you can require that the NETCONF server enforce certain behaviors during the NETCONF session by configuring the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level. When you configure the **rfc-compliant** statement, the NETCONF server explicitly declares the NETCONF namespace in its replies and qualifies all NETCONF tags with the **nc** prefix. Also, **<get>** and **<get-config>** operations that return no configuration data do not include an empty **<configuration>** element in RPC replies.

[See [Configuring RFC-Compliant NETCONF Sessions.](#)]

- **New YANG features including configuration hierarchy must constraints published in YANG and a new module that defines Junos OS YANG extensions (PTX Series)**—Starting in Junos OS Release 15.1, the Juniper Networks **configuration** YANG module includes configuration constraints published using either the YANG **must** statement or the Junos OS YANG extension **junos:must**. Constraints that cannot be mapped directly to YANG's **must** statement, which include expressions containing special keywords or symbols such as **all**, **any**, **unique**, **\$**, **\_**, and wildcard characters, are published using **junos:must**.

The new **junos-extension** module contains definitions for Junos OS YANG extensions including the **must** and **must-message** keywords. The **junos-extension** module is bound to the namespace URI **http://yang.juniper.net/yang/1.1/je** and uses the prefix **junos**. You can download Juniper Networks YANG modules from the website, or you can generate the modules by using the **show system schema** operational mode command on the local device.

[See [Using Juniper Networks YANG Modules.](#)]

## MPLS

---

- **New command to display the MPLS label availability in RPD (PTX Series)**—Starting with Junos OS Release 15.1, a new show command, **show mpls label usage**, is introduced to display the available label space resource in RPD and also the applications that use the label space in RPD. Using this command, the administrator can monitor the available labels in each label space and the applications that are using the labels.

[See [show mpls label usage](#).]

## Routing Protocols

---

- **BGP PIC for inet (PTX Series)**—Beginning with Junos OS Release 15.1, BGP Prefix Independent Convergence (PIC), which was initially supported for Layer 3 VPN routers, is extended to BGP with multiple routes in the global tables such as inet and inet6 unicast, and inet and inet6 labeled unicast. When the BGP PIC feature is enabled on a router, BGP installs to the Packet Forwarding Engine the second best path in addition to the calculated best path to a destination. When an IGP loses reachability to a prefix, the router uses this backup path to reduce traffic loss until the global convergence through the BGP is resolved, thereby drastically reducing the outage duration.

[See [Use Case for BGP PIC for Inet](#).]

- **Multi-instance support for RSVP-TE (PTX Series)**—Beginning with Junos OS Release 15.1, multi-instance support is extended to the existing MPLS RSVP-Traffic Engineering (TE) functionality. This support is available only for a virtual router instance type. A router can create and participate in multiple independent TE topology partitions, which allows each partitioned TE domain to scale independently.

Multi-instance support is also extended for LDP over RSVP tunneling for a virtual router routing instance. This supports splitting of a single routing and MPLS domain into multiple domains so that each domain can be scaled independently. BGP labeled unicast can be used to stitch these domains for service FECs. Each domain uses intra-domain LDP over RSVP LSP for MPLS forwarding.

[See [Tunneling LDP LSPs in RSVP LSPs Overview](#).]

- **Selection of backup LFA for OSPF routing protocol (PTX Series)**—Starting with Junos OS Release 15.1, the default loop-free alternative (LFA) selection algorithm or criteria can be overridden with an LFA policy. These policies are configured per destination per primary next-hop interface or per destination. These backup policies enforce LFA selection based on admin-group, srlg, node, bandwidth, protection-type, and metric attributes of the backup path. During backup shortest-path-first (SPF) computation, each attribute (both node and link) of the backup path, stored per backup next hop, is accumulated by IGP. For the routes created internally by IGP, the attribute set of every backup path is evaluated against the policy configured per destination per prefix primary next-hop interface. The first or the best backup path is selected and installed as the backup next hop in the routing table.

[See [Example-configuring-backup-selection-policy-for-ospf-protocol](#).]

- **Remote LFA support for LDP in OSPF (PTX Series)**—Beginning with Junos OS Release 15.1, you can configure a remote loop-free alternate (LFA) to extend the backup provided

by the LFA in an OSPF network. This feature is useful especially for Layer 1 metro-rings where the remote LFA is not directly connected to the PLR. The existing LDP implemented for the MPLS tunnel setup can be reused for the protection of OSPF networks and subsequent LDP destinations thereby eliminating the need for RSVP-TE backup tunnels for backup coverage.

[See [Example-configuring-remote-lfa-over-ldp-tunnels-in-ospf-networks](#).]

## User Interface and Configuration

- **Support for displaying configuration differences in XML tag format (PTX Series)**—Starting with Junos OS Release 15.1, you can use the **show compare | display xml** command to compare the candidate configuration with the current committed configuration and display the differences between the two configurations in XML tag format.

[See [Understanding the show | compare | display xml Command Output](#).]

- **Configuring chassis ambient temperature to optimize the power consumption of FPCs (PTX5000)**—Starting with Junos OS Release 15.1, the power management feature of the PTX5000 is enhanced to manage the power supplied to the FPCs by configuring the ambient temperature of the chassis. You can set the ambient temperature of the chassis at 25° C or 40° C. On system initialization, the power manager reads the ambient temperature and allocates power to the FPCs according to the power budget policy at that temperature. If any FPC consumes more power than the configured value for more than 3 minutes, the **PWR Range Overshoot** alarm is raised for that FPC, and the power manager overrides the configured ambient temperature setting of that FPC and resets its ambient temperature to the next higher level and reallocates power according to the new temperature setting. All the overshooting FPCs remain in the dynamic ambient temperature mode until the next reboot, or until you override it with a CLI command. The power manager then resets the power budget of the FRUs, including the overshooting FPCs, according to the configured ambient temperature setting.

To configure the ambient temperature, include the **set chassis ambient-temperature** statement at the **[edit]** hierarchy level.



**NOTE:** If ambient temperature is not configured, then default ambient temperature is set as 55° C.

[See [Chassis Ambient-Temperature](#).]

## VPNs

---

- **Segmented inter-area P2MP LSP (PTX Series)**—Starting with Junos OS Release 15.1, P2MP LSPs can be segmented at the area boundary. A segmented P2MP LSP consists of an ingress area segment (ingress PE router or ASBR), backbone area segment (transit ABR), and egress area segment (egress PE routers or ASBRs). Each of the intra-area segments can be carried over provider tunnels such as P2MP RSVP-TE LSP, P2MP mLDP LSP, or ingress replication. Segmentation of inter-area P2MP LSP occurs when the S-PMSI auto-discovery routes are advertised, which triggers the inclusion of a new BGP extended community or inter-area P2MP segmented next-hop extended community in the ingress PE router or ASBR, transit ABR, and egress PE routers or ASBRs.

[See [Example: Configuring Segmented Inter-Area P2MP LSP](#).]

### Related Documentation

- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Known Issues on page 122](#)
- [Resolved Issues on page 124](#)
- [Documentation Updates on page 127](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

## Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 15.1R2 for the PTX Series.

- [High Availability \(HA\) and Resiliency on page 120](#)
- [Junos OS XML API and Scripting on page 121](#)
- [Routing Protocols on page 121](#)
- [User Interface and Configuration on page 121](#)

### High Availability (HA) and Resiliency

---

- **A check option is added for command request chassis routing-engine master (all platforms)**—Starting in Junos OS Release 15.1, a **check** option available with the **switch**, **release**, and **acquire** options checks the GRES status of the standby Routing Engine before toggling mastership. The **force** option is also removed from all platforms.

[See [request chassis routing-engine master](#).]

- **GRES readiness is part of show system switchover output (PTX Series)**—Starting in Junos OS Release 15.1, switchover readiness status is reported as part of the output for operational mode command **show system switchover**.



## Junos OS XML API and Scripting

- **Escaping of special XML characters required for request\_login (PTX Series)**—Beginning with Junos OS Release 15.1R2, you must escape any special characters in the username and password elements of a **request\_login** XML RPC request. The following five symbols are considered special characters: greater than (>), less than (<), single quote ('), double quote ("), and ampersand (&). Both entity references and character references are acceptable escape sequence formats. For example, **&amp;** and **&#38;** are valid representations of an ampersand. Previously no escaping of these characters was required.

## Routing Protocols

- **New IS-IS adjacency holddown CLI command (PTX Series)**—Beginning with Junos OS Release 15.1, a new operational command **show isis adjacency holddown** is introduced to display the adjacency holddown status. This command is useful to verify whether the adjacency holddown is enabled and facilitates troubleshooting when there are adjacency issues due to IS-IS adjacency holddown.

[See [show isis adjacency holddown](#).]

- **Configure and establish targeted sessions with third-party controllers using LDP targeted neighbor (PTX Series)**—Starting with Junos OS Release 15.1, you can configure LDP targeted neighbor to third-party controllers for applications such as route recorder that wants to learn label-FEC bindings of an LSR. LDP targeted neighbor helps to establish a targeted session with controllers for a variety of applications.

## User Interface and Configuration

- **Changed available REST interface cipher suites when Junos OS is in FIPS mode (PTX Series)**—Starting with Junos OS Release 15.1, when Junos OS is in FIPS mode, you can only configure cipher suites with a FIPS-compliant hash algorithm for the REST interface to the device. To configure a cipher suite, specify the **cipher-list** statement at the **[edit system services rest]** hierarchy level.

[See [cipher-list \(REST API\)](#).]

- **New flag to control errors when executing multiple RPCs through a REST interface (PTX Series)**—Starting with Junos OS Release 15.1, you can stop on an error when executing multiple RPCs through a REST interface by specifying the **stop-on-error** flag in the HTTP POST method.

[See [Submitting a POST Request to the REST API](#).]

### Related Documentation

- [New and Changed Features on page 114](#)
- [Known Behavior on page 122](#)
- [Known Issues on page 122](#)
- [Resolved Issues on page 124](#)
- [Documentation Updates on page 127](#)

- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

## Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R2 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.

- [System Logging on page 122](#)

---

### System Logging

- **Text string deprecated in syslog messages that are converted to SNMP traps (PTX Series)**—In the syslog messages that are converted to SNMP traps for event policies, the "trap sent successfully" text string is deprecated.

#### Related Documentation

- [New and Changed Features on page 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Issues on page 122](#)
- [Documentation Updates on page 127](#)
- [Resolved Issues on page 124](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

## Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R2 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 122](#)
- [Interfaces and Chassis on page 123](#)
- [Network Management and Monitoring on page 123](#)
- [Routing Protocols on page 123](#)
- [Software Installation and Upgrade on page 124](#)

---

### General Routing

- The PTX Series does not support the queuing PICs, but by default Junos OS will program the chassis scheduler map which will generate the following logs: "fpc2 COS(cos\_chassis\_scheduler\_pre\_add\_action:2140): chassis scheduler ipc received for

non qplic 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." [PR910985](#)

- When a switchover is done from one Routing Engine to the other, in graceful-switchover redundancy mode, there is a brief period early in the transition of the SIB to online state, during which unsolicited (not corresponding to an attempt by the CPU to access the SIB via PCIe) errors are received at the downstream PCIe port on the CB to the SIB. The fix is to mute the generation of such errors during this brief period of the switchover. [PR1068237](#)
- Starting with Junos OS Release 14.1, Entropy Label Capability is enabled by default on all Juniper Networks [PTX] systems. On PTX Series transit LSRs that carry LSPs with Entropy Label Capability, packet loss can be observed due to data errors when one or more labeled route entries are not properly removed from the hash table (ie. following LSP optimization or MBB event) because the 'stale' entries are pointing to corrupted route memory. As a result, when the MPLS label that's associated with the 'stale' entry is re-used, data errors are seen for packets using the corresponding label. [PR1100637](#)

### Interfaces and Chassis

- On PTX Series routers, TX optical threshold value is shown incorrect for the interfaces in the PIC P1-PTX-2-100G-WDM. This PR will fix only the TX power issue reported in the 2x100G DWDM OTN PIC. [PR1084963](#)
- On PTX 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](#)

### Network Management and Monitoring

- If Routing Engine protocol mastership is not established and a daemon like mib2d tries to register with shm-rtddb for ifState updates, it may not receive updates. Due to this, a recent fix was introduced to delay the above registration until Routing Engine's Protocol Mastership is resolved. As a side effect of this fix - we see this core. In this case SNMP Requests have landed on the mib2d, before it has connected to shm-rtddb and initialized its interface database. As a fix - (To Avoid SNMP Requests landing on mib2d before Routing Engine-Mastership is resolved) we have delayed the MIB registration as well. Hence after Routing Engine bootup, once Protocol Mastership is resolved - mib2d will connect to shm-rtddb and then register its MIBs with snmpd. So no snmp requests will be received in mib2d until mastership is resolved. [PR1114001](#)

### Routing Protocols

- On shmlog unsupported platforms, the following message might be seen after a configuration change: PTX-re0 rpd[42030] shmlog not initialized for PIM - not provisioned in platform manifest file The message does not indicate an error, it just indicates that shmlog is not supported on the PTX Series platform. The severity of the log has been reduced to INFO. [PR1065055](#)
- In IS-IS environment, MPLS LSPs are established, when IS-IS traceoptions flag "general" is activated, the LSP convergence time is increased. [PR1090752](#)

## Software Installation and Upgrade

---

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

### Related Documentation

- [New and Changed Features on page 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Resolved Issues on page 124](#)
- [Documentation Updates on page 127](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

## 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.

- [Forwarding and Sampling on page 125](#)
- [General Routing on page 125](#)
- [Interfaces and Chassis on page 126](#)
- [MPLS on page 127](#)
- [Network Management and Monitoring on page 127](#)
- [Routing Protocols on page 127](#)

## 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 may never come out of that loop which may result in high CPU usage (up to 90% sometimes). Because, sampled is not able to consume any states (such as route updates, interface updates) generated by kernel and this results in memory exhaustion and finally resulting in the router not making any updates and forcing a router reboot. [PR1092684](#)

## General Routing

- On PTX Series platform, when performing scaling (for example, polling 768 IFDs via SNMP with max of 92 PPS and with all 8 FPCs online) SNMP polling on the device, due to the large number of messages between Routing Engine and Packet Forwarding Engine, PFEMAN (Packet Forwarding Engine manager) errors might be seen on the router, which may cause high SNMP response time and CPU spike (for example, increase 8% when executing the "show" command) as well. [PR1078003](#)
- PTX Series Packet Forwarding Engine does not support L3VPN VRF and we can assign only loopback (lo0) interface to VRF as management VRF, so returning commit error by applying non-loopback interface under vrf instance is correct. # commit check [edit routing-instances l3vpn interface] 'et-8/0/0.0' RT Instance: Only loopback interface is supported under vrf routing instances. error: configuration check-out failed In 14.1, we see the same commit error when a non-loopback interface is configured under vrf instance on PTX3K, while in 14.2, commit goes through without any error. Without the commit error, customer may encounter packet discard issue when mistakenly configuring L3VPN PE with PTX3K. This is a PTX3K specific issue with 14.2. If we try 14.2 on PTX5K, we will see the commit error. [PR1078960](#)
- Tunable SFP+ optics will not be supported on P1-PTX-24-10G-W-SFPP PIC in Junos OS 15.1R1 release. On Tunable Optics in this PIC, with 15.1R1, the wavelength will not be configurable and the tunable parameters will not be correctly displayed in the CLI. [PR1081992](#)
- The FPC on PTX Series router might crash and reboot when the Packet Forwarding Engine is handling a fatal error, when the error happened, "TQCHIP0: Fatal error pqt\_min\_free\_cnt is zero" log message will be seen. [PR1084259](#)
- Due to a change in an existing PR, group names in the configuration must be a string of alphanumericals, dashes or underscores. There is not workaround other than following the group name instructions. [PR1087051](#)
- In Dual Routing Engine systems when both Routing Engines reboot and after coming up if the mastership is not established or takes time to establish, mib2d may start and exit 4 times in quick succession. Hence it will not be running. As a workaround, it can be simply started again once Routing Engine-Mastership is established. This is a race condition and hence may not be seen always. [PR1087428](#)
- On PTX Series platforms, some non-fatal interrupts (for example, CM cache or AQD interrupts) are logged as fatal interrupts. The following log messages will be shown

on CM parity interrupt: fpc0 TQCHIP 0: CM parity Fatal interrupt, Interrupt status: 0x10  
fpc0 CMSNG: Fatal ASIC error, chip TQ fpc0 TQCHIP 0: CM cache parity Fatal interrupt  
has occurred 181 time(s) in 180010 msec TQCHIP 0: CM cache parity Fatal interrupt  
has occurred 181 time(s) in 180005 msec [PR1089955](#)

- On Junos OS Release 15.1R1, when the multicast next-hop is changed, the grafting and pruning operations take more time than before. [PR1090608](#)
- When the PTX Series only has bits-a and bits-b as configured clock sources (and there is no interface on FPC configured as clock source), and it is losing signal from both of bits-a and bits-b simultaneously, clock sync state will go to FREERUN mode immediately, this is unexpected behavior. After the fix of this PR, clock sync state will stay HOLDOVER, then will go to FREERUN mode after the timeout. [PR1099516](#)
- On PTX Series platform, when yanking out FPC or SIB ungracefully (for example, pulling the line card out of the chassis unintentionally when the line card is carrying the traffic), there might be small probability that it can impact any of the FPCs with Grant Scheduler (GS) and Request Table (RT) fatal interrupt occurred. [PR1105079](#)

---

## Interfaces and Chassis

- If we load 15.1 Junos jinstall/jinstall64 image on PTX Series and if we have CFM configured over AE interfaces, the FPC might crash. [PR1085952](#)
- 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](#)
- After removing a child link from AE bundle, in the output of "show interface <AE> detail", the packets count on the remaining child link spikes, then if add back the previous child link, the count recover to normal. [PR1091425](#)
- On PTX Series platform, if the configurations that have per-unit-scheduler configured on the interface, but without proper class-of-service configuration for the same interface, due to lack of commit check, the device control daemon (dcd) may fail to return "commit error" and pass the configuration. Following is an example, user@re0# set interfaces et-0/0/1 per-unit-scheduler vlan-tagging unit 0 <<<<< The configuration for interface et-0/0/1 user@re0# commit check error: per-unit-scheduler is configured but class-of-service is blank <<<<< This is correct behavior error: configuration check-out failed <<<<< .. user@re0# set class-of-service forwarding-classes queue 7 q7 <<<<< user@re0# commit check configuration check succeeds <<<<< This is wrong behavior because et-0/0/1 does not have class-of-service configuration \* If reboot this router after committing, the administrator cannot access without console because the router cannot read this configuration. When deleting the above configuration after rebooting, telnet etc could be used. [PR1097829](#)

## MPLS

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

## Network Management and Monitoring

- Due to inappropriate cleanup in async library, disable multiple interfaces while SNMP is polling interface oids might cause mid2d process to crash. [PR1097165](#)

## Routing Protocols

- On PTX Series platform with transit BGP-LU chained composite next-hop configured, when advertising LDP routes via BGP labeled unicast (BGP-LU), if the LDP LSP itself is tunneled over an RSVP LSP, the rpd process might crash. Notes: The "set routing-options forwarding-table chained-composite-next-hop transit labeled-bgp" is enabled by default on PTX Series. [PR1065107](#)

### Related Documentation

- [New and Changed Features on page 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Documentation Updates on page 127](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

## Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R2 documentation for the PTX Series.

- [High Availability Feature Guide on page 127](#)
- [IPv6 Neighbor Discovery Feature Guide on page 128](#)

### High Availability Feature Guide

- The following information belongs in the "Nonstop Active Routing Concepts" topic:  
If you have NSR configured, it is never valid to issue the **restart routing** command in any form on the NSR master Routing Engine. Doing so results in a loss of protocol adjacencies and neighbors and a drop in traffic.
- The following information belongs in the "Configuring Nonstop Active Routing" topic:  
If the routing protocol process (rpd) on the NSR master Routing Engine crashes, the master Routing Engine simply restarts rpd (with no Routing Engine switchover), which impacts routing protocol adjacencies and neighbors and results in traffic loss. To prevent this negative impact on traffic flow, configure the **switchover-on-routing-crash**

statement at the **[edit system]** hierarchy level. This configuration forces an NSR Routing Engine switchover if rpd on the master Routing Engine crashes.

---

### IPv6 Neighbor Discovery Feature Guide

---

- The “NDP Cache Protection Overview,” “Configuring NDP Cache Protection,” “Example: Configuring NDP Cache Protection to Prevent Denial-of-Service Attacks,” and “nd-system-cache-limit” topics failed to include the EX Series, M Series, PTX Series, and T Series as supported platforms. These platforms, as well as the MX series, are all supported.

#### Related Documentation

- [New and Changed Features on page 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Known Issues on page 122](#)
- [Resolved Issues on page 124](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)
- [Product Compatibility on page 131](#)

## 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 128](#)
- [Upgrading a Router with Redundant Routing Engines on page 129](#)
- [Basic Procedure for Upgrading to Release 15.1 on page 129](#)

---

### 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 15.1

---

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 15.1R1 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 15.1R1 jinstall package, you cannot issue the request system software rollback command to return to the previously installed software. Instead you must issue the request system software add validate command and specify the jinstall package that corresponds to the previously installed software.

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

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

```
user@host> request system software add validate reboot source/jinstall-15.1  
R21-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-15.1
R21-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 15.1 **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 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Known Issues on page 122](#)
- [Resolved Issues on page 124](#)
- [Documentation Updates on page 127](#)
- [Product Compatibility on page 131](#)

## Product Compatibility

- [Hardware Compatibility on page 132](#)

## 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 114](#)
- [Changes in Behavior and Syntax on page 120](#)
- [Known Behavior on page 122](#)
- [Known Issues on page 122](#)
- [Resolved Issues on page 124](#)
- [Documentation Updates on page 127](#)
- [Migration, Upgrade, and Downgrade Instructions on page 128](#)

## 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

18 February 2016—Revision 6, Junos OS Release 15.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

14 January 2016—Revision 5, Junos OS Release 15.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

20 November 2015—Revision 4, Junos OS Release 15.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

9 November 2015—Revision 3, Junos OS Release 15.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

3 November 2015—Revision 2, Junos OS Release 15.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

26 October 2015—Revision 1, Junos OS Release 15.1R2— EX Series, M Series, MX Series, PTX Series, and T Series.

25 August 2015—Revision 6, Junos OS Release 15.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

23 July 2015—Revision 5, Junos OS Release 15.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

2 July 2015—Revision 4, Junos OS Release 15.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

19 June 2015—Revision 3, Junos OS Release 15.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

12 June 2015—Revision 2, Junos OS Release 15.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

5 June 2015—Revision 1, Junos OS Release 15.1R1— EX Series, 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.