

Junos[®] OS 15.1 Release Notes

INSIDE THIS RELEASE

- Supported on EX Series, M Series, MX Series, PTX Series, and T Series

NEW SOFTWARE FEATURES

- Media Access Control Security (MACsec) (EX9204, EX9208, EX9214)
- Junos OS SNMP scripts to support custom MIBs (M Series, MX Series, and T Series)
- MX Series Virtual Chassis support for MX2010 and MX2020 member routers (MX Series with MPCs/MICs)
- Captive portal content delivery (HTTP redirect) supported on MS-MICs, MS-MPCs, and the Routing Engine (MX Series)
- IGMP Snooping on Pseudowires (MX Series)
- Sender-Based RPF and hot-root standby for ingress replication (IR) provider tunnels (MX Series)
- Fast-failover according to flow rate (MX Series with MPCs)
- Extended MPC support for per-unit schedulers (MX Series)
- Support for dynamic power management (MX Series)
- Support for flexible queuing on non-HQoS MPCs (MX Series)
- Leveraging DPCs for EVPN deployment (MX Series with DPCs)
- Multi-instance support for RSVP-TE (M Series, MX Series, and T Series)
- Entropy label support for BGP-LU (MX Series with MPCs, and T Series with HC-FPC)
- Pro-active loss and delay measurement (MX Series routers with MPCs and MICs only)
- Segmented inter-area P2MP LSP (M Series, MX Series, T Series)
- Support for NSR, IGP-FA, and static route on container LSPs (M Series, MX Series, and T Series)
- BGP Prefix Independent Convergence for inet (MX Series routers with MPCs)
- Configuring per-prefix LFA and node to link protection fallback for OSPF (M Series, MX Series and T Series)
- Remote LFA support for LDP in OSPF (MX Series)
- Configuring per-interface NDP cache protection (MX Series)
- Media Access Control Security (MACsec) (MX240, MX480, and MX960)

NEW HARDWARE FEATURES

- 2-port 100-Gigabit Metro DWDM OTN PIC (PTX Series)
- New MPC variants that support higher scale and bandwidth (MX Series)



Junos[®] OS Release 15.1R1 for the EX Series, M Series, MX Series, PTX Series, and T Series

25 August 2015

Contents

Introduction	5
Junos OS Release Notes for EX Series Switches	5
New and Changed Features	5
Interfaces and Chassis	6
Junos OS XML API and Scripting	7
Management	7
MPLS	8
Port Security	8
Software Installation and Upgrade	9
Spanning-Tree Protocols	9
Changes in Behavior and Syntax	9
Known Behavior	10
Authentication and Access Control	10
J-Web	10
Port Security	11
Spanning-Tree Protocols	11
Virtual Chassis	11
Known Issues	11
Infrastructure	12
J-Web	12
Software-Defined Networking (SDN)	12
Software Installation and Upgrade	12
Documentation Updates	13
Migration, Upgrade, and Downgrade Instructions	13
Upgrade and Downgrade Support Policy for Junos OS Releases	14

Product Compatibility	14
Hardware Compatibility	14
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D Universal Edge Routers, and T Series Core Routers	16
New and Changed Features	16
Hardware	17
Bridging and Learning	17
Class of Service (CoS)	18
High Availability (HA) and Resiliency	19
Interfaces and Chassis	21
IPv6	24
Junos OS XML API and Scripting	24
Layer 2 Features	24
Management	25
Multicast	26
MPLS	27
Network Management and Monitoring	28
Routing Policy and Firewall Filters	29
Routing Protocols	30
Services Applications	32
Subscriber Management and Services (MX Series)	35
Software Installation and Upgrade	40
User Interface and Configuration	40
VPNs	40
Changes in Behavior and Syntax	42
General Routing	43
High Availability (HA) and Resiliency	43
MPLS	44
Routing Protocols	44
Routing Policy and Firewall Filters	47
Security	47
Services Applications	47
Subscriber Management and Services (MX Series)	49
User Interface and Configuration	53
VPNs	53
Known Behavior	53
Subscriber Management and Services (MX Series)	54
Known Issues	54
Class of Service (CoS)	55
Forwarding and Sampling	55
General Routing	55
Interfaces and Chassis	57
J-Web	58
Layer 2 Features	58
MPLS	58
Network Management and Monitoring	58
Platform and Infrastructure	59
Routing Policy and Firewall Filters	60
Routing Protocols	60

Services Applications	60
Software-Defined Networking (SDN)	61
User Interface and Configuration	61
VPNs	61
Documentation Updates	62
High Availability Feature Guide	62
IPv6 Neighbor Discovery Feature Guide for Routing Devices	63
Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices	63
MPLS Applications Feature Guide for Routing Devices	63
Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices	63
Subscriber Management Provisioning Guide	63
Migration, Upgrade, and Downgrade Instructions	64
Basic Procedure for Upgrading to Release 15.1	65
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)	67
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)	68
Upgrade and Downgrade Support Policy for Junos OS Releases	70
Upgrading a Router with Redundant Routing Engines	70
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1	70
Upgrading the Software for a Routing Matrix	72
Upgrading Using Unified ISSU	73
Downgrading from Release 15.1	73
Product Compatibility	74
Hardware Compatibility	74
Junos OS Release Notes for PTX Series Packet Transport Routers	75
New and Changed Features	75
Hardware	75
Interfaces and Chassis	76
IPv6	77
Junos OS XML API and Scripting	77
Management	78
MPLS	79
Routing Protocols	79
User Interface and Configuration	80
VPNs	81
Changes in Behavior and Syntax	81
High Availability (HA) and Resiliency	81
Routing Protocols	82
User Interface and Configuration	82
Known Behavior	82
System Logging	83
Known Issues	83
General Routing	83
Interfaces and Chassis	84
Documentation Updates	84
High Availability Feature Guide	84

Migration, Upgrade, and Downgrade Instructions	85
Upgrading Using Unified ISSU	85
Upgrading a Router with Redundant Routing Engines	85
Basic Procedure for Upgrading to Release 15.1R1	86
Product Compatibility	88
Hardware Compatibility	89
Third-Party Components	90
Finding More Information	90
Documentation Feedback	90
Requesting Technical Support	91
Self-Help Online Tools and Resources	91
Opening a Case with JTAC	91
Revision History	92

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, and T Series, and Junos Fusion.

These release notes accompany Junos OS Release 15.1R1 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 15.1R1 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/>.



NOTE: On EX9200 switches, OpenFlow and OVSDDB are not supported in Junos OS Release 15.1R1.8.

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 11](#)
- [Documentation Updates on page 13](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)
- [Product Compatibility on page 14](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R1 for the EX Series.



NOTE: The following EX Series platforms are supported in Release 15.1R1: 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 6](#)
 - [Junos OS XML API and Scripting on page 7](#)
 - [Management on page 7](#)
 - [MPLS on page 8](#)
 - [Port Security on page 8](#)
 - [Software Installation and Upgrade on page 9](#)
 - [Spanning-Tree Protocols on page 9](#)

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 Communication 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 limit 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 within 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 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 changes more than the configured number of times within one second, the changes to MAC addresses are dropped, logged or ignored, or the interface is shut down.

[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 different output 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 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 11](#)
- [Resolved Issues](#)
- [Documentation Updates on page 13](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)
- [Product Compatibility on page 14](#)

Changes in Behavior and Syntax

There are no changes in default behavior and syntax in Junos OS Release 15.1R1 for EX Series switches.

Related Documentation

- [New and Changed Features on page 5](#)
- [Known Behavior on page 10](#)

- [Known Issues on page 11](#)
- *Resolved Issues*
- [Documentation Updates on page 13](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)
- [Product Compatibility on page 14](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R1 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 name (filter name plus term name plus counter name) that is longer than 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 total length of those three names is added to the length of the interface name and the MAC address, the total length 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 when 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 deleted the VoIP configuration, deleted the current data VLAN membership, and configured 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. If at the same time the client clears the IP address on the same interface, the clearance of the IP address on interface message, generated by the kernel (event), is processed at the client and triggers the DHCP client state machine, which leads to acquisition of 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 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 **bpdud-block-on-edge** statement in a VSTP configuration on the switch. [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 5](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Issues on page 11](#)
- [Resolved Issues](#)
- [Documentation Updates on page 13](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)
- [Product Compatibility on page 14](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R1 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.

- [Infrastructure on page 12](#)
- [J-Web on page 12](#)

- [Software-Defined Networking \(SDN\) on page 12](#)
- [Software Installation and Upgrade on page 12](#)

Infrastructure

- On EX2200 switches, syslog 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 packets)**. The correct log message would be: **PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 10.0.1.114 10.0.0.7 0 0 (1 packets)**. [PR898175](#)

J-Web

- If you access the J-Web interface using 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, 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, on the Route Information page (Monitor > Routing > Route Information), the Next Hop column in 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 corresponding IP address of the next-hop interface. [PR603669](#)
- 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-Defined Networking (SDN)

- On EX9200 switches, OpenFlow is not supported in Junos OS Release 15.1R1.8.
- On EX9200 switches, OVSD is not supported in Junos OS Release 15.1R1.8.

Software Installation and Upgrade

- On a mixed EX4200 and EX4500 Virtual Chassis or on an EX3300 Virtual Chassis, during a Nonstop Software Upgrade (NSSU) to Release 15.1R1, you might see duplicate packets. [PR1062944](#)
- Substantial traffic losses might occur when you run an NSSU upgrade on EX4200 and EX4500 Virtual Chassis, EX6200 and EX8200 switches, or EX8200 Virtual Chassis. [PR1062960](#)

-
- On an EX8200 Virtual Chassis, a Nonstop Software Upgrade (NSSU) to Release 15.1R1 might fail after the image is pushed to the backup Routing Engine, and a vmcore might be created. [PR1075232](#)
 - On EX9200 switches, an ISSU upgrade to Release 15.1R1 might fail during the switchover from master Routing Engine to standby Routing Engine. [PR1088827](#)
 - On EX9200 switches, ISSU is not working properly for upgrading to Junos OS Release 15.1R1. Junos Space is triggering the upgrades and the upgrades are failing. [PR1091610](#)
 - 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 on the device image, and select **ex-92xx** in **Modify device image**. [PR1090863](#)

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Resolved Issues](#)
- [Documentation Updates on page 13](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)
- [Product Compatibility on page 14](#)

Documentation Updates

There are no errata or changes in Junos OS Release 15.1R1 for the EX Series switches documentation.

Related Documentation

- [New and Changed Features on page 5](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 11](#)
- [Resolved Issues](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)
- [Product Compatibility on page 14](#)

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 14](#)

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 earlier or later. 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 5](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 11](#)
- [Resolved Issues](#)
- [Documentation Updates on page 13](#)
- [Product Compatibility on page 14](#)

Product Compatibility

- [Hardware Compatibility on page 14](#)

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 5](#)
- [Changes in Behavior and Syntax on page 9](#)
- [Known Behavior on page 10](#)
- [Known Issues on page 11](#)
- *Resolved Issues*
- [Documentation Updates on page 13](#)
- [Migration, Upgrade, and Downgrade Instructions on page 13](#)

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 15.1R1 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.



NOTE: On MX Series routers, OpenFlow and OVSDDB are not supported in Junos OS Release 15.1R1.8.

- [New and Changed Features on page 16](#)
- [Changes in Behavior and Syntax on page 42](#)
- [Known Behavior on page 53](#)
- [Known Issues on page 54](#)
- [Documentation Updates on page 62](#)
- [Migration, Upgrade, and Downgrade Instructions on page 64](#)
- [Product Compatibility on page 74](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R1 for the M Series, MX Series, and T Series.

- [Hardware on page 17](#)
- [Bridging and Learning on page 17](#)
- [Class of Service \(CoS\) on page 18](#)
- [High Availability \(HA\) and Resiliency on page 19](#)
- [Interfaces and Chassis on page 21](#)
- [IPv6 on page 24](#)
- [Junos OS XML API and Scripting on page 24](#)
- [Layer 2 Features on page 24](#)
- [Management on page 25](#)
- [Multicast on page 26](#)

- [MPLS on page 27](#)
- [Network Management and Monitoring on page 28](#)
- [Routing Policy and Firewall Filters on page 29](#)
- [Routing Protocols on page 30](#)
- [Services Applications on page 32](#)
- [Subscriber Management and Services \(MX Series\) on page 35](#)
- [Software Installation and Upgrade on page 40](#)
- [User Interface and Configuration on page 40](#)
- [VPNs on page 40](#)

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.

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)**—Junos OS Release 15.1 or later enables you to 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](#).]

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**—As of 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)**—As of Junos OS Release 15.1, there is a feedback mechanism that 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)**—As of 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**—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

- **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.

- **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.

- **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.

- **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](#).]

- **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 \(MS-DPC\)](#).]

- **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 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.

- **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](#).]

- **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 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.
- **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](#).]

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, 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](#).]

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.](#)]

Multicast

- **Latency fairness optimized multicast (MX Series)**—Starting with Junos OS Release 15.1, 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. These options are supported only on platforms with the **enhanced-ip** mode enabled.

[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 (IR) provider tunnels (MX Series routers with MPCs running in "enhanced-ip" mode)**—Junos OS Release 15.1 and later releases add support for sender-based RPF and hot-root standby to IR 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 IR 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 IR to distinguish the sending PE router and the VRF. Note that when IR is used as the selective provider tunnel, IR 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. IR 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 \(MBGP MVPN\)](#).]

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 through a backup AC, while in the Centralized protector model, the PLR or the protector has an MPLS tunnel to the backup PE.

[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 (FA), 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-FA 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](#).]

Network Management and Monitoring

- **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.

Routing Policy and Firewall Filters

- **Support for consistent load balancing for ECMP groups (MX Series routers with MPCs)**—Effective 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, new forwarding class accounting statistics can be enabled 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.](#)]

- **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.](#)]

- **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.](#)]

- **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](#).]

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 Jflow) 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 Jflow) 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.

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.

Software Installation and Upgrade

- **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](#).]

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 15.1R1.

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 42](#)
- [Known Behavior on page 53](#)
- [Known Issues on page 54](#)
- [Documentation Updates on page 62](#)
- [Migration, Upgrade, and Downgrade Instructions on page 64](#)
- [Product Compatibility on page 74](#)

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 15.1R1 for the M Series, MX Series, and T Series.

- [General Routing on page 43](#)
- [High Availability \(HA\) and Resiliency on page 43](#)
- [MPLS on page 44](#)
- [Routing Protocols on page 44](#)
- [Routing Policy and Firewall Filters on page 47](#)
- [Security on page 47](#)
- [Services Applications on page 47](#)
- [Subscriber Management and Services \(MX Series\) on page 49](#)
- [User Interface and Configuration on page 53](#)
- [VPNs on page 53](#)

General Routing

- **commit synchronize statement is not allowed in batch mode**—When user attempts **commit atomic** in configure batch mode, a warning is shown to the user: "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 **set system commit synchronize** statement followed by **commit**.

High Availability (HA) and Resiliency

- **VRRP adjusted priority can go to zero (M Series, MX Series, and T Series)**—As of 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**—As of Junos OS Release 15.1, there is a **check** option available with the **switch**, **release**, and **acquire** options that 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)**—As of 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](#).]

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](#).]

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)**— 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 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. Note that 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.

Routing Policy and Firewall Filters

- **Command completion for the [show firewall prefix-action-stats filter *filter-name* prefix-action] hierarchy on all compatible platforms**—Prior to 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.

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** *mapping-lifetime-max* statement at the **[edit services pcp]** hierarchy level is modified to be 0–4294667, instead of the previous range 0–2147483647.

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 **anccp** 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-ift-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-ift-scaling** option.

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 \(REST API\)](#).]

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 16](#)
- [Known Behavior on page 53](#)
- [Known Issues on page 54](#)
- [Documentation Updates on page 62](#)
- [Migration, Upgrade, and Downgrade Instructions on page 64](#)
- [Product Compatibility on page 74](#)

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R1 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.

- [Subscriber Management and Services \(MX Series\) on page 54](#)

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.

Related Documentation

- [New and Changed Features on page 16](#)
- [Changes in Behavior and Syntax on page 42](#)
- [Known Issues on page 54](#)
- [Documentation Updates on page 62](#)
- [Migration, Upgrade, and Downgrade Instructions on page 64](#)
- [Product Compatibility on page 74](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R1 for the M Series, MX Series and T Series.

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

- [Class of Service \(CoS\) on page 55](#)
- [Forwarding and Sampling on page 55](#)
- [General Routing on page 55](#)
- [Interfaces and Chassis on page 57](#)
- [J-Web on page 58](#)
- [Layer 2 Features on page 58](#)
- [MPLS on page 58](#)
- [Network Management and Monitoring on page 58](#)
- [Platform and Infrastructure on page 59](#)
- [Routing Policy and Firewall Filters on page 60](#)
- [Routing Protocols on page 60](#)
- [Services Applications on page 60](#)
- [Software-Defined Networking \(SDN\) on page 61](#)

-
- [User Interface and Configuration on page 61](#)
 - [VPNs on page 61](#)

[Class of Service \(CoS\)](#)

- 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](#)
- If Layer2 egress sampling and egress dot1p rewrite are both set on an interface, the sampled packet will not have the rewritten dot1p bits. [PR1081203](#)

[Forwarding and Sampling](#)

- The L2ALD daemon may generate a core file whenever a configured mesh group or routing instance type change is followed by a logical system delete. If this does occur the presence of a core file is an indication that this problem has been encountered. The system will function normally after the daemon restarts. [PR914404](#)

[General Routing](#)

- On MX Series routers equipped with MPCs, ping fails with packet size greater than 4000 bytes if max-queues-per-interface configured as 4 explicitly. [PR902525](#)
- Traceroute or SSD crash seen when as-number-lookup option is used when executing traceroute. [PR928769](#)
- When BCM0 interface goes down, Routing Engine should switch over on M320. [PR949517](#)
- A PE device running EVPN IRB with an IGP configured in a VRF associated with the EVPN instance will be unable to establish an IGP adjacency with a CE device attached to a remote PE. The IGP instance running in the VRF on the PE may be able to discover the IGP instance running on the remote CE through broadcast or multicast traffic, but will be unable to send unicast traffic directly to the remote CE. [PR977945](#)
- In point-to-point (P2P) SONET/SDH interface environment, there is a destination route with this interface as next-hop. When this interface is disabled, the destination route is still kept in the forwarding table and might cause ping fails with "Can't assign requested address" error. [PR984623](#)
- If vrf-target of an EVPN routing-instance that has interfaces with ESI configured, is changed, per-ES AD type 1 routes are not exchanged between Provider Edge routers in that EVPN instance. This could affect functions like mass withdrawal and L2 load balancing to multihomed CE. [PR990931](#)
- This PR is implementing traceoptions debug enhancements to detect route-record corruption events. The route-record traceoptions debug will be enabled as follows:
----- user@router> edit Entering configuration mode [edit]
user@router# set routing-options traceoptions flag route-record [edit] user@router#
commit ----- [PR1015820](#)

- When trying to view the address pooling and/or Endpoint independent mappings from a particular private or a public IP address, all the mappings will be displayed. [PR1019739](#)
- In the scenario where router acts as both egress LSP for core network and BRAS for subscribers, RSVP-TE sends PathErr to ingress router due to matching to subscriber interfaces incorrectly when checking the explicit route object (ERO), if subscribers are associated with same lo0 address as used by RSVP LSP egress address. [PR1031513](#)
- With an unrecognized or unsupported Control Board (CB), mismatch link speed might be seen between fabric and FPCs, which results in FPCs CRC/destination errors and fabric planes offline. Second issue is in a race condition. Fabric Manager (FM) might process the stale destination disable event but the error is cleared instead, which results in the unnecessary FPC offline and not allowing Fabric Hardening action to trigger and recover. [PR1031561](#)
- When the CPU usage is very high (e.g. 100%) on Routing Engine, the MS-MIC might get stuck due to kernel deadlock, which triggers the card to crash and generate a core file. [PR1038026](#)
- Queue stats on LSQ interfaces are not properly cleaned up when queuing is enabled on the IFD and the queues are hosted at IFD level. This happens with a subsequent deletion and creation of LSQ interface (not always though). [PR1044340](#)
- 1 to 3 seconds of traffic loss is seen during local repair of LSPs using fast-reroute(not link protection). The link used by LSPs is brought down using laser off from remote end. [PR1048109](#)
- Routing protocol process (rpd) might crash with core-dump due to memory leak. [PR1052614](#)
- The MS-MPC does not support clock synchronization as it has no clocking capable interfaces. Upon receipt of these clock synchronization messages the router will log the following: Jan 20 17:33:14.032 2015 ROUTER_RE0 : %PFE-3: fpc1 gencfg no msg handlers for gencfg msg command 34 Jan 20 17:35:18.388 2015 ROUTER_RE0 /kernel: %KERN-1-GENCFG: op 34 (CLKSYNC blob) failed; err 7 (Doesn't Exist) This PR will allow the MS-MPC to ignore these messages and prevent the logs from being generated. [PR1062132](#)
- You might see higher baseline CPU utilization and periodic CPU spikes seen on MPC6E Cards as compared to 16x10GE MPC Cards. On MPC6E, we have low priority threads that monitor various things in the background on a periodic basis such as voltage, temperature, stats counters, and hardware status. When the system is idle, these threads are allowed to take more of the load, and that is why we are seeing higher baseline CPU/CPU spikes. This does not prevent other higher priority threads from running when they have to, as these are noncritical activities being done in the background and hence is a nonimpacting issue. [PR1071408](#)
- 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. root@sdn1-qfx24q-a> show ovsdb netconf transactions Txn ID Logical-switch Port VLAN ID 1 vlan100 root@sdn1-qfx24q-a> clear ovsdb netconf transactions. [PR1072730](#)

-
- Traffic throughput when flowing from MPCs and MICs to ROHS2-Compliant MPCs is less when compared to ROHS2-Compliant MPCs to MPCs and MICs. [PR1076009](#)
 - If RTSP (Real Time Streaming Protocol) ALG has been configured, MS-MIC might crash with core-dump in scaled application layer traffic environment. [PR1076573](#)
 - Multiple negative tests such a restarting routing or chassis-control may cause the router to reboot. [PR1077428](#)
 - From Junos OS Release 14.1R1, if the hidden knob "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](#)
 - 'show interfaces queue <ifl>' stats are not correct with RLSQ warm-standby mode. Issue seen on MPCs and MICs. [PR1082417](#)

Interfaces and Chassis

- Packet Forwarding Engine continues to forward traffic to DHCP client on a demux interface when ae0 interface is down. In this scenario the AE interface bundle has five members and configured with minimum link value of 4. When two members are down, the ae0 interface also goes down, but Packet Forwarding Engine continues to forward traffic on other members for the demux interface. [PR836846](#)
- Time taken to reboot T Series boxes has gone up. T Series(Standalone) 14.2 - 3 minutes 39 seconds; 15.1 - 4 minutes 18 seconds (Difference - 40s) TX Matrix (Multichassis) 14.2 - 5 minutes 17 seconds; 15.1 - 7 minutes 18 seconds (Difference - 2 minutes). [PR1049869](#)
- There is a mismatch in mac statistics, few frames go unaccounted. This is a day 1 issue with the software fetching of mac statistics. The snap and clear bits were set together on pm3393 chip driver software, so it used to happen that even before the copy of stats to shadow registers happened, clear was happening which used to go unaccounted. Now rollover mechanism has been implemented and tested for two continuous days and everything is fine. [PR1056232](#)
- When "set chassis lcc 0 offline" is used on SCC and committed, the configuration gets synced on LCC. However when "delete chassis lcc 0 offline" is used on SCC, we need to do commit two times on SCC in order to sync the configuration on LCC being brought online. [PR1058994](#)
- 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. [PR1060659](#)
- Link Up/Down SNMP traps for AE member links might not be generated, but the SNMP traps for the AE bundle works well. [PR1067011](#)
- MX104 show chassis fpc errors are not applicable to MX80/ MX104; hence this command is being removed. [PR1071553](#)

- After toggling the enhanced-ip mode and rebooting the router, FPC alarms may still show for an FPC which is powered down after the reboot. [PR1082851](#)
- MAC accounting statistics cannot be displayed for an aggregated Ethernet interface in routed mode, when the same MAC is learned on multiple child links of the bundle. This issue is seen only on the MX80 and MX104. The CLI command "show interface <name> mac-database" times out. [PR1082862](#)

J-Web

- On HTTPS service, J-Web is not launching the chassis viewer page at Internet Explorer 7. [PR819717](#)
- On configure->clitools->point and click->system->advanced->deletion of saved core context on "No" option is not happening at J-Web. [PR888714](#)
- Basic value entry format error check is not present in Configure-->Security-->IPv6 Firewall Filters, but the same is present in IPv4 Firewall Filters. But it will throw error when try to commit the wrong format data entered. [PR1009173](#)

Layer 2 Features

- With Dynamic Host Configuration Protocol (DHCP) maintain subscriber feature enabled, when the subscriber's incoming interface index is changed, for example, the interfaces go away and come back after changing the MTU configuration of interface, the existing subscribers may get dropped and new subscribers fail in connection. [PR1059999](#)
- After Routing Engine switchover as part of GRES, sometimes we can see a momentarily flood of data frames due to delay in re-convergence of xSTP (VSTP, MSTP, RSTP) topology. [PR1064225](#)

MPLS

- RSVP graceful restart does not function for LSPs that have a forwarding adjacency (FA) label-switched path (LSP) as a next hop. [PR60256](#)
- For point-to-multipoint LSPs configured for VPLS, the "ping mpls" command reports 100 percent packet loss even though the VPLS connection is active. [PR287990](#)
- In scenario of egress-protection using stub-alias advertise mode where Point of Local Repair (PLR) use 'dynamic-rsvp-lsp' in LDP link protection, if protected PE get isolated, unexpected packet drops will be observed. [PR1030815](#)

Network Management and Monitoring

- SNMP mibs jnxFWCounterByteCount, jnxFWCounterDisplayFilterName, jnxFWCounterDisplayName, jnxFWCounterDisplayType may be missing from jnxFirewallCounterTable when "show snmp mib walk jnxFirewallCounterTable" is executed. [PR1040043](#)
- In some race conditions with firewall filters change, it is possible that the mib2d process receives a new MX Series filter ADD event before it learns about a non-MX Series filter

DELETE event for the same filter index. The mib2d process will crash due to this. [PR1057373](#)

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](#)
- Wrong source IP is used when responding to traceroute in L3VPN setup. This is not an indication of traffic taking wrong link. [PR883701](#)
- The overhead values need to be represented with 8 bits to cover the range "-120..124", but the microcode is only using the last 7 bits. [PR1020446](#)
- On MX Series based platform, with igmp-snooping enabled and a multicast route with integrated routing and bridging (IRB) as a downstream interface, a multicast composite nexthop is created with a list of L3 and corresponding L2 nexthops. In a rare corner case, the corresponding L2 nexthop to the L3 IRB nexthop is a DISCARD nexthop and will cause the FPC to crash. [PR1026124](#)
- A Packet Forwarding Engine memory leak is seen when multicast receivers are connected in a bridge domain where IGMP snooping is enabled and IGMP messages are exchanged between the multicast receivers and the Layer 3 IRB (integrated routing and bridging) interface. [PR1027473](#)
- 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](#)
- An FPC with interfaces configured as part of an aggregated Ethernet bundle may core and reboot when the shared-bandwidth-policer is configured as part of the firewall policer. [PR1069763](#)
- If with about 1M routes on MX Series router, there might be more than 1 second (about 1.3s) packets dark window during unified ISSU. [PR1070217](#)
- High volume of VC internal chassis to chassis TCP control flow can impact VC stability and responsiveness to external protocol events. The solution is to move this TCP control traffic onto a parallel VC host path punt queue separate from VC control. DDOS Virtual Chassis protocol statistics provide a mechanism to verify the VC punt queue receive activity. [PR1074760](#)
- When a MX chassis network-services is "enhanced-ip" and an AE with "family bridge" configuration is first committed, there is a possibility that an incorrect forwarding path may be installed causing traffic loss. [PR1081999](#)
- On MX Series platform, the "RPF-loose-mode-discard" feature is not working when configured within a Virtual Router routing instance. The feature is working for the routing instance only when configured in the main instance. [PR1084715](#)

Routing Policy and Firewall Filters

- Executing CLI command "show route resolution" and stopping the command output before reaching the end of the database, the rpd process might crash when executing the same command again. [PR1023682](#)
- When there are more than 1000 routes, the "show multicast route extensive" command , takes more than 2 minutes to display the output. [PR1084983](#)

Routing Protocols

- It is necessary that the MSDP peer local-address matches the PIM RP address on routers that are RP. MSDP RPF check might fail in rare cases when both these addresses are not equal. [PR35806](#)
- When you configure damping globally and use the import policy to prevent damping for specific routes, and a peer sends a new route that has the local interface address as the next hop, the route is added to the routing table with default damping parameters, even though the import policy has a nondefault setting. As a result, damping settings do not change appropriately when the route attributes change. [PR51975](#)
- Continuous soft core-dump may be observed due to bgp-path-selection code. RPD forks a child and the child asserts to produce a core-dump. The problem is with route-ordering. And it is auto-corrected after collecting this soft-assert-coredump, without any impact to traffic/service. [PR815146](#)
- In rare cases, rpd may write a core file with signature "rt_notbest_sanity: Path selection failure on ..." The core is 'soft', which means there should be no impact to traffic or routing protocols. [PR946415](#)
- 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 right upstream. Nbr results in join not propagated to the upstream nbr . Shows command for pim join upstream nbr "unknown" . [PR1069896](#)

Services Applications

- When you specify a standard application at the [edit security idp idp-policy <policy-name> rulebase-ips rule <rule-name> match application] hierarchy level, IDP does not detect the attack on the nonstandard port (for example, junos:ftp on port 85). Whether it is a custom or predefined application, the application name does not matter. IDP simply looks at the protocol and port from the application definition. Only when traffic matches the protocol and port does IDP try to match or detect against the associated attack. [PR477748](#)
- 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](#)
- jl2tpd crash in L2tp::sendProxyLcpAuthData. [PR1082673](#)

Software-Defined Networking (SDN)

- On MX Series routers, OpenFlow is not supported in Junos OS Release 15.1R1.8.
- On MX Series routers, OVSDDB is not supported in Junos OS Release 15.1R1.8.

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

VPNs

- When you modify the frame-relay-tcc statement at the [edit interfaces interface-name unit logical-unit-number] hierarchy level of a Layer 2 VPN, the connection for the second logical interface might not come up. As a workaround, restart the chassis process (chassisd) or reboot the router. [PR32763](#)
- It is planned 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 will 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](#)
- On a dual Routing Engine, 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](#)

- Related Documentation**
- [New and Changed Features on page 16](#)
 - [Changes in Behavior and Syntax on page 42](#)
 - [Known Behavior on page 53](#)
 - [Documentation Updates on page 62](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 64](#)
 - [Product Compatibility on page 74](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R1 documentation for the M Series, MX Series, and T Series.

- [High Availability Feature Guide on page 62](#)
- [IPv6 Neighbor Discovery Feature Guide for Routing Devices on page 63](#)
- [Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices on page 63](#)
- [MPLS Applications Feature Guide for Routing Devices on page 63](#)
- [Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices on page 63](#)
- [Subscriber Management Provisioning Guide on page 63](#)

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 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](#).]

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.

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.

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 Juniper Networks recommends 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.

Subscriber Management Provisioning 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 one minute between subscriber login minus the 10-second timeout).

Related Documentation

- [New and Changed Features on page 16](#)
- [Changes in Behavior and Syntax on page 42](#)
- [Known Behavior on page 53](#)
- [Known Issues on page 54](#)
- [Migration, Upgrade, and Downgrade Instructions on page 64](#)
- [Product Compatibility on page 74](#)

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 65](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) on page 67](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 68](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 70](#)
- [Upgrading a Router with Redundant Routing Engines on page 70](#)
- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 70](#)
- [Upgrading the Software for a Routing Matrix on page 72](#)
- [Upgrading Using Unified ISSU on page 73](#)
- [Downgrading from Release 15.1 on page 73](#)

[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 **jbundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).



.....

NOTE: Before upgrading, back up the file system and the currently active Junos OS configuration so that you can recover to a known, stable environment in case the upgrade is unsuccessful. Issue the following command:

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

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the `juniper.conf` and `ssh` files) might be removed. To preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the [Junos OS Administration Library for Routing Devices](#).

.....

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

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



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

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

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



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

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-32-15.1R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot  
source/junos-install-mx-x86-64-15.1R1.9-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 Junos OS Release 15.1 M Series, MX Series, and T Series routers.

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

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

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



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

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

```
user@host> request system software add validate reboot  
source/jinstall-15.1R1.9-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-15.1R1.9-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 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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

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

Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the *Installation and Upgrade Guide*.

Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1

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

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

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



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

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

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



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

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

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

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

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

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



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

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

Upgrading the Software for a Routing Matrix

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

- A minimum of free disk space and DRAM on each Routing Engine. The software upgrade will fail on any Routing Engine without the required amount of free disk space and DRAM. To determine the amount of disk space currently available on all Routing Engines of the routing matrix, use the CLI **show system storage** command. To determine the amount of DRAM currently available on all the Routing Engines in the routing matrix, use the CLI **show chassis routing-engine** command.
- The master Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC are all re0 or are all re1.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC are all re1 or are all re0.
- All master Routing Engines in all routers run the same version of software. This is necessary for the routing matrix to operate.
- All master and backup Routing Engines run the same version of software before beginning the upgrade procedure. Different versions of Junos OS can have incompatible message formats especially if you turn on GRES. Because the steps in the process include changing mastership, running the same version of software is recommended.
- For a routing matrix with a TX Matrix router, the same Routing Engine model is used within a TX Matrix router (SCC) and within a T640 router (LCC) of a routing matrix. For example, a routing matrix with an SCC using two RE-A-2000s and an LCC using two RE-1600s is supported. However, an SCC or an LCC with two different Routing Engine models is not supported. We suggest that all Routing Engines be the same model throughout all routers in the routing matrix. To determine the Routing Engine type, use the CLI **show chassis hardware | match routing** command.
- For a routing matrix with a TX Matrix Plus router, the SFC contains two model RE-DUO-C2600-16G Routing Engines, and each LCC contains two model RE-DUO-C1800-8G or RE-DUO-C1800-16G Routing Engines.



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

To upgrade the software for a routing matrix:

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

For the detailed procedure, see the [Routing Matrix with a TX Matrix Router Deployment Guide](#) or the [Routing Matrix with a TX Matrix Plus Router Deployment Guide](#).

Upgrading Using Unified ISSU



CAUTION: This release introduces some behavior changes to the unified in-service software upgrade (ISSU) functionality for M, MX, 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 16](#)
- [Changes in Behavior and Syntax on page 42](#)
- [Known Behavior on page 53](#)
- [Known Issues on page 54](#)
- [Documentation Updates on page 62](#)
- [Product Compatibility on page 74](#)

Product Compatibility

- [Hardware Compatibility on page 74](#)

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 15.1R1 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 75](#)
- [Changes in Behavior and Syntax on page 81](#)
- [Known Behavior on page 82](#)
- [Known Issues on page 83](#)
- [Documentation Updates on page 84](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)
- [Product Compatibility on page 88](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R1 for the PTX Series.

- [Hardware on page 75](#)
- [Interfaces and Chassis on page 76](#)
- [IPv6 on page 77](#)
- [Junos OS XML API and Scripting on page 77](#)
- [Management on page 78](#)
- [MPLS on page 79](#)
- [Routing Protocols on page 79](#)
- [User Interface and Configuration on page 80](#)
- [VPNs on page 81](#)

Hardware

- **2-port 100-Gigabit Metro DWDM OTN PIC (PTX Series)**—Starting in Junos OS Release 15.1, the 2-port 100-Gigabit Metro DWDM OTN PIC (PTX-2-100G-WDM-M) is supported on PTX3000 and PTX5000 routers.

The PIC supports:

- Metro applications
- Transparent transport of two 100-Gigabit Ethernet signals with OTU4 framing
- ITU-standard OTN performance monitoring and alarm management
- Dual polarization quadrature phase shift keying (DP-QPSK) modulation and soft-decision forward error correction (SD-FEC)

[See [PTX Series Interface Module Reference](#) and [100-Gigabit Ethernet OTN Options Configuration Overview](#).]

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-SFP) 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-SFP 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)**—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 81](#)
- [Known Behavior on page 82](#)
- [Known Issues on page 83](#)
- [Documentation Updates on page 84](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)
- [Product Compatibility on page 88](#)

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 15.1R1 for the PTX Series.

- [High Availability \(HA\) and Resiliency on page 81](#)
- [Routing Protocols on page 82](#)
- [User Interface and Configuration on page 82](#)

High Availability (HA) and Resiliency

- **A check option is added for command request chassis routing-engine master (all platforms)**—As of Junos OS Release 15.1, there is a **check** option available with the **switch**, **release**, and **acquire** options that 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)**—As of Junos OS Release 15.1, switchover readiness status is reported as part of the output for operational mode command **show system switchover**.

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](#).]

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 75](#)
- [Known Behavior on page 82](#)
- [Known Issues on page 83](#)
- [Documentation Updates on page 84](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)
- [Product Compatibility on page 88](#)

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R1 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 83](#)

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 75](#)
- [Changes in Behavior and Syntax on page 81](#)
- [Known Issues on page 83](#)
- [Documentation Updates on page 84](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)
- [Product Compatibility on page 88](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R1 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 83](#)
- [Interfaces and Chassis on page 84](#)

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 qpvc ifd et-2/1/3 with index 131 /kernel: GENCFG: op 8 (COS BLOB) failed; err 5 (Invalid)Fix: Adding check to stop sending chassis scheduler map on PTX platform." [PR910985](#)
- PTX Series Packet Forwarding Engine does not support L3VPN VRF. 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 Release 14.1, we see the same commit error when a non-loopback interface is configured under vrf instance on the PTX3000, while in Release 14.2, commit goes through without any error. Without the commit error, customer might encounter packet discard issue when mistakenly configuring L3VPN PE with the PTX3000. This is a PTX3000 specific issue with Release 14.2. If we try Release 14.2 on PTX5000, we see the commit error. [PR1078960](#)

- If we load jinstall/jinstall64 image on PTX Series router and if we have CFM configured over AE interfaces, this issue will be seen. [PR1085952](#)
- Tunable SFP+ optics will not be supported on P1-PTX-24-10G-W-SFPP in 15.1R1 release. On Tunable Optics in this PIC, with Release 15.1R1, the Wavelength will not be configurable and the tunable parameters will not be correctly displayed in the CLI. [PR1081992](#)

Interfaces and Chassis

- With the current design, when we add/delete filters on aggregated interfaces - bwy halp receives both parent and child iff changes for output filters but in case of input halp only receives the parent iff change message. We go through the child list in parent iff and add all the child bind points. This creates a problem when some of the child ifls go down/up in the bundle. halp might not receive messages in this case. This problem was resolved as part of PR 863789. Fix for this PR calls newly added functions that add/delete bind points in halp accordingly. But there is a memory corruption on this path. Because of corruption in this path, when FPC is rebooted, this code patch creates some corrupted bind points in halp. This does not lead to the FPC crash. But, when the bind points are deleted from CLI, this causes a crash in the IFF message processing path while deleting the bind points. [PR1066795](#)

Related Documentation

- [New and Changed Features on page 75](#)
- [Changes in Behavior and Syntax on page 81](#)
- [Known Behavior on page 82](#)
- [Documentation Updates on page 84](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)
- [Product Compatibility on page 88](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R1 documentation for the PTX Series.

- [High Availability Feature Guide on page 84](#)

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.

Related Documentation

- [New and Changed Features on page 75](#)
- [Changes in Behavior and Syntax on page 81](#)
- [Known Behavior on page 82](#)
- [Known Issues on page 83](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)
- [Product Compatibility on page 88](#)

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 85](#)
- [Upgrading a Router with Redundant Routing Engines on page 85](#)
- [Basic Procedure for Upgrading to Release 15.1R1 on page 86](#)

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 15.1R1

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  
R11-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-15.1  
R11-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 75](#)
- [Changes in Behavior and Syntax on page 81](#)
- [Known Behavior on page 82](#)
- [Known Issues on page 83](#)
- [Documentation Updates on page 84](#)
- [Product Compatibility on page 88](#)

Product Compatibility

- [Hardware Compatibility on page 89](#)

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 75](#)
- [Changes in Behavior and Syntax on page 81](#)
- [Known Behavior on page 82](#)
- [Known Issues on page 83](#)
- [Documentation Updates on page 84](#)
- [Migration, Upgrade, and Downgrade Instructions on page 85](#)

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 <https://www.juniper.net/cgi-bin/docbugreport/>.
- E-mail—Send your comments to 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

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 © 2015, 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.