

Release Notes: Junos[®] OS Release 15.1R3 for the ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series

10 June 2016

Contents

Introduction	7
Junos OS Release Notes for ACX Series	7
New and Changed Features	7
Hardware	8
Class of Service	8
Firewall Filters	9
Interfaces and Chassis	10
Installation	16
Layer 2 Features	16
Management	21
Routing	22
Security	22
Subscriber Access Management	22
Timing and Synchronization	22
Changes in Default Behavior and Syntax	23
Interfaces and Chassis	24
Known Behavior	24
Known Issues	25
Class of Service	25
Firewall Filters	26
Interfaces and Chassis	28
Integrated Routing and Bridging	30
Layer 2 Services	31
MPLS Applications	32
Network Management	32
Statistics	33

Timing and Synchronization	33
Resolved Issues	33
Resolved Issues	33
Documentation Updates	34
Migration, Upgrade, and Downgrade Instructions	34
Upgrade and Downgrade Support Policy for Junos OS Releases	34
Product Compatibility	35
Hardware Compatibility	35
Junos OS Release Notes for EX Series Switches	36
New and Changed Features	36
Hardware	37
Authentication and Access Control	38
Interfaces and Chassis	39
Junos OS XML API and Scripting	40
Management	40
MPLS	41
Port Security	41
Software Installation and Upgrade	42
Spanning-Tree Protocols	42
Changes in Behavior and Syntax	43
Dynamic Host Configuration Protocol	43
Known Behavior	43
Authentication and Access Control	43
Interfaces and Chassis	44
J-Web	44
Network Management and Monitoring	45
Port Security	45
Software Installation and Upgrade	45
Spanning-Tree Protocols	46
Virtual Chassis	46
Known Issues	46
High Availability (HA) and Resiliency	46
Infrastructure	47
Port Security	47
Software Installation and Upgrade	47
Resolved Issues	47
Resolved Issues: Release 15.1R3	48
Resolved Issues: Release 15.1R2	53
Documentation Updates	56
Migration, Upgrade, and Downgrade Instructions	56
Upgrade and Downgrade Support Policy for Junos OS Releases	57
Product Compatibility	57
Hardware Compatibility	57
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 3D	
Universal Edge Routers, and T Series Core Routers	59
New and Changed Features	59
Hardware	60
Bridging and Learning	61
Class of Service (CoS)	61

High Availability (HA) and Resiliency	62
Interfaces and Chassis	64
IPv6	69
Junos OS XML API and Scripting	69
Layer 2 Features	69
Management	71
MPLS	71
Multicast	73
Network Management and Monitoring	75
Routing Policy and Firewall Filters	76
Routing Protocols	77
Services Applications	80
Software Defined Networking	84
Software Installation and Upgrade	84
Subscriber Management and Services (MX Series)	85
User Interface and Configuration	98
VPNs	98
Changes in Behavior and Syntax	100
Authentication, Authorization and Accounting	101
Class of Service (CoS)	101
General Routing	101
High Availability (HA) and Resiliency	102
Junos XML API and Scripting	104
Layer 2 VPNs	104
MPLS	104
Multicast	104
Network Management and Monitoring	104
Routing Policy and Firewall Filters	106
Routing Protocols	106
Security	109
Services Applications	109
Subscriber Management and Services (MX Series)	111
System Logging	117
System Management	124
User Interface and Configuration	124
Virtual Chassis	125
VPNs	125
Known Behavior	125
Hardware	125
MPLS	126
Subscriber Management and Services (MX Series)	126
System Logging	127
Known Issues	128
Class of Service	128
Forwarding and Sampling	128
General Routing	129
Infrastructure	132
Interfaces and Chassis	132
J-Web	134

Layer 2 Features	134
MPLS	134
Network Management and Monitoring	135
Platform and Infrastructure	136
Routing Protocols	138
Services Applications	139
Subscriber Management and Services	140
User Interface and Configuration	141
VPNs	141
Resolved Issues	142
Resolved Issues: 15.1R3	143
High Availability (HA) and Resiliency	156
Infrastructure	157
Interfaces and Chassis	158
Layer 2 Features	162
MPLS	163
Network Management and Monitoring	165
Platform and Infrastructure	165
Routing Protocols	170
Routing Policy and Firewall Filters	172
Services Applications	173
Software Installation and Upgrade	174
Subscriber Management and Services	174
User Interface and Configuration	177
VPNs	177
Resolved Issues: 15.1R2	178
Documentation Updates	200
Adaptive Services Interfaces Feature Guide for Routing Devices	200
Broadband Subscriber VLANs and Interfaces Feature Guide	201
High Availability Feature Guide	201
IPv6 Neighbor Discovery Feature Guide for Routing Devices	202
Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices	202
MPLS Applications Feature Guide for Routing Devices	203
Overview for Routing Devices	204
Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices	204
Standards Reference	204
Security Services Administration Guide for Routing Devices	204
Subscriber Management Provisioning Guide	204
User Access and Authentication Guide for Routing Devices	205
VPNs Library for Routing Devices	205
Migration, Upgrade, and Downgrade Instructions	205
Basic Procedure for Upgrading to Release 15.1	206
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)	208
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)	209
Upgrade and Downgrade Support Policy for Junos OS Releases	211
Upgrading a Router with Redundant Routing Engines	211

Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1	211
Upgrading the Software for a Routing Matrix	213
Upgrading Using Unified ISSU	214
Downgrading from Release 15.1	214
Product Compatibility	215
Hardware Compatibility	215
Junos OS Release Notes for PTX Series Packet Transport Routers	216
New and Changed Features	216
High Availability and Resiliency (HA)	217
Interfaces and Chassis	217
IPv6	218
Junos OS XML API and Scripting	218
Management	219
MPLS	220
Routing Protocols	220
User Interface and Configuration	221
VPNs	222
Changes in Behavior and Syntax	222
High Availability (HA) and Resiliency	222
Junos OS XML API and Scripting	223
Network Management and Monitoring	223
Routing Protocols	223
User Interface and Configuration	223
Known Behavior	224
System Logging	224
Known Issues	225
General Routing	225
MPLS	226
Routing Protocols	226
Resolved Issues	226
Resolved Issues: 15.1R3	226
Resolved Issues: 15.1R2	229
Documentation Updates	232
High Availability Feature Guide	232
IPv6 Neighbor Discovery Feature Guide	232
Migration, Upgrade, and Downgrade Instructions	233
Upgrading Using Unified ISSU	233
Upgrading a Router with Redundant Routing Engines	233
Basic Procedure for Upgrading to Release 15.1	233
Product Compatibility	236
Hardware Compatibility	237
Junos OS Release Notes for the QFX Series	238
New and Changed Features	238
Management	238
Network Management and Monitoring	240
Spanning-Tree Protocols	240

User Interface and Configuration	240
Changes in Behavior and Syntax	241
Routing Protocols	241
Known Behavior	241
Virtual Chassis	241
Known Issues	242
Firewall Filters	242
High Availability	242
Infrastructure	242
Interfaces and Chassis	242
Routing Policy	243
Resolved Issues	243
Resolved Issues: Release 15.1R3	243
Documentation Updates	249
Migration, Upgrade, and Downgrade Instructions	250
Upgrading Software on QFX3500, QFX3600, and QFX5100 Standalone Switches	250
Performing an In-Service Software Upgrade (ISSU) on the QFX5100 Switch	251
Product Compatibility	254
Hardware Compatibility	254
Third-Party Components	255
Finding More Information	255
Documentation Feedback	255
Requesting Technical Support	256
Self-Help Online Tools and Resources	256
Opening a Case with JTAC	256
Revision History	257

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, EX Series, M Series, MX Series, PTX Series, QFabric systems, QFX Series, SRX Series, T Series, and Junos Fusion.

These release notes accompany Junos OS Release 15.1R3 for the ACX Series, EX Series, M Series, MX Series, PTX Series, QFX 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 ACX Series

These release notes accompany Junos OS Release 15.1R3 for the ACX 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/beta/junos/>.

- [New and Changed Features on page 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Documentation Updates on page 34](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

New and Changed Features

This section describes the features and enhancements in Junos OS Release 15.1R3 for ACX Series Universal Access Routers.

- [Hardware on page 8](#)
- [Class of Service on page 8](#)
- [Firewall Filters on page 9](#)
- [Interfaces and Chassis on page 10](#)
- [Installation on page 16](#)
- [Layer 2 Features on page 16](#)
- [Management on page 21](#)
- [Routing on page 22](#)
- [Security on page 22](#)
- [Subscriber Access Management on page 22](#)
- [Timing and Synchronization on page 22](#)

Hardware

- **ACX Series Universal Access Router**—Starting with Junos OS Release 15.1R3, Junos OS supports the following Juniper Networks ACX Series Universal Access Routers:
 - [ACX1000 and ACX1100 Universal Access Router](#)
 - [ACX2000 and ACX2100 Universal Access Router](#)
 - [ACX2200 Universal Access Router](#)
 - [ACX4000 Universal Access Router](#)

These routers enable a wide range of business and residential applications and services, including microwave cell site aggregation, MSO mobile backhaul service cell site deployment, and service provider or operator cell site deployment.

Class of Service

- **Class of service for PPP and MLPPP interfaces (ACX Series)**—Junos OS for ACX Series Universal Access Routers support class-of-service (CoS) functionalities on PPP and MLPPP interfaces. Up to four forwarding classes and four queues are supported per logical interface for PPP and MLPPP packets.

The following restrictions apply when you configure CoS on PPP and MLPPP interfaces on ACX Series routers:

- For interfaces with PPP encapsulation, you can configure interfaces to support only the IPv4, Internet Protocol Control Protocol (IPCP), PPP Challenge Handshake Authentication Protocol (CHAP), and Password Authentication Protocol (PAP) applications.
 - Drop timeout is not supported.
 - Loss of traffic occurs during a change of scheduling configuration; you cannot modify scheduling attributes instantaneously.
 - Buffer size is calculated in terms of number of packets, with 256 bytes considered as the average packet size.
 - Only two loss priority levels, namely low and high, are supported.
- **Support for MLPPP encapsulation (ACX Series)**—You configure multilink bundles as logical units or channels on the link services interface `lsq-0/0/0`. With MLPPP, multilink bundles are configured as logical units on `lsq-0/0/0`—for example, `lsq-0/0/0.0` and `lsq-0/0/0.1`. After creating multilink bundles, you add constituent links to the bundle.

MLPPP is supported on ACX1000, ACX2000, and ACX2100 routers, and with Channelized OC3/STM1 (Multi-Rate) MICs with SFP and 16-port Channelized E1/T1 Circuit Emulation MIC on ACX4000 routers. With multilink PPP bundles, you can use the PPP Challenge Handshake Authentication Protocol (CHAP) and Password Authentication Protocol (PAP) for secure transmission over the PPP interfaces.

To configure MLPPP encapsulation, include the **encapsulation multilink-ppp** statement at the `[edit interfaces lsq-fpc/pic/port unit logical-unit-number]` hierarchy level. To

aggregate T1 links into a an MLPPP bundle, include the **bundle** statement at the **[edit interfaces t1-fpc/pic/port unit *logical-unit-number* family mlppp]** hierarchy level.

- **Support for configuring the shared buffer size (ACX Series)**—Junos OS for ACX Series Universal Access Routers enable you to control the amount of shared packet buffer a given queue can consume. Using this feature, you can ensure that important queues have a higher chance of using the shared buffers than by not so important queues. To achieve this, you can configure lower values for **shared-buffer maximum** CLI statement for the not so important queues, and higher values for the **shared-buffer maximum** CLI statement for the important queues.

You can explicitly configure the **shared-buffer maximum** CLI statement at the **[edit class-of-service]** hierarchy level.



NOTE: The default value for **shared-buffer maximum** is 66%.

Firewall Filters

- **Support for hierarchical policers (ACX Series)**—On ACX Series routers, two-level ingress hierarchical policing is supported. With single-level policers, you cannot administer the method using which the committed information rate (CIR) and the excess information rate (EIR) values specified in the bandwidth profile are shared across different flows. For example, in a certain network deployment, you might want an equal or even distribution of CIR across the individual flows. In such a scenario, you cannot accomplish this requirement using single-level policers and need to configure aggregate or hierarchical policers.

Aggregate policers operate in peak, guarantee, and hybrid modes. You can configure an aggregate policer by including the **aggregate-policer *aggregate-policer-name*** statement at the **[edit firewall policer *policer-name* if-exceeding]** hierarchy level. You can specify the mode of the aggregate policer by including the **aggregate-sharing-mode [guarantee | peak | hybrid]** statement at the **[edit firewall policer *policer-name* if-exceeding aggregate-policer *aggregate-policer-name*]** hierarchy level.

- **Enhancement to support additional firewall filter match capabilities (ACX Series)**—Starting in Release 12.3X54, Junos OS for ACX Series router supports additional match capabilities at the **[edit firewall family ccc filter]** and **[edit firewall family inet filter]** hierarchy levels.

The existing firewall do not support Layer 2, Layer 3, and Layer 4 fields at the **[edit firewall family ccc filter]** hierarchy level. With additional matching fields, ACX Series routers support all the available Layer 2, Layer 3, and Layer 4 fields on the user-to-network interface side (ethernet-ccc/vlan-ccc).

At the **[edit firewall family inet filter]** hierarchy level, the **fragment-flags** match field has been removed to accommodate the following Layer 2 and Layer 3 fields:

Table 1: Fields added to [edit firewall family inet filter] hierarchy level

Field	Description
first-fragment	Matches if packet is the first fragment
is-fragment	Matches if packet is a fragment

The scale for **inet** and **ccc** in the firewall family filter has been reduced from 250 hardware entries to 122 hardware entries.

Interfaces and Chassis

- **Support for Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (ACX4000)**—The ACX4000 Universal Access Routers support the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (model number ACX-MIC-4COC3-1COC12CE).

The key features supported are:

- Structure-Agnostic TDM over Packet (SAToP)
- Pseudowire Emulation Edge to Edge (PWE3) control word for use over an MPLS packet-switched network (PSN)
- **Support for 6-port Gigabit Ethernet Copper/SFP MIC (ACX4000)**—The ACX4000 Universal Access Routers support the 6-port Gigabit Ethernet Copper/SFP MIC. The 6-port Gigabit Ethernet Copper/SFP MIC features six tri-speed (10/100/1000 Mbps) Ethernet ports. Each port can be configured to operate in either RJ45 or SFP mode and can support PoE.
- **Support for chassis management (ACX4000)**—The ACX4000 Universal Access Routers support the following CLI operational mode commands:

Show commands:

- **show chassis alarms**
- **show chassis craft-interface**
- **show chassis environment**
- **show chassis environment pem**
- **show chassis fan**
- **show chassis firmware**
- **show chassis fpc *pic-status***
- **show chassis hardware (clei-models | detail | extensive | models)**
- **show chassis mac-addresses**
- **show chassis pic fpc-slot *fpc-slot pic-slot pic slot***
- **show chassis routing-engine**

Restart command:

- **restart chassis-control** (*gracefully | immediately | soft*)

Request commands:

- **request chassis feb restart slot slot-number**
- **request chassis mic mic-slot *mic-slot* fpc-slot *fpc-slot* (offline | online)**
- **request chassis pic offline fpc-slot *fpc-slot* pic-slot *pic-slot***
- **User-defined alarms (ACX Series)**—On an ACX Series router, the alarm contact port (labeled ALARM) provides four user-defined input ports and two user-defined output ports. Whenever a system condition occurs—such as a rise in temperature, and depending on the configuration, the input or output port is activated.

To view the alarm relay information, issue the **show chassis craft-interface** command from the Junos OS command-line interface.

- **Support for Ethernet synthetic loss measurement (ACX Series)**—You can trigger on-demand and proactive Operations, Administration, and Maintenance (OAM) for measurement of statistical counter values corresponding to ingress and egress synthetic frames. Frame loss is calculated using synthetic frames instead of data traffic. These counters maintain a count of transmitted and received synthetic frames and frame loss between a pair of maintenance association end points (MEPs).

The Junos OS implementation of Ethernet synthetic loss measurement (ETH-SLM) is fully compliant with the ITU-T Recommendation Y.1731. Junos OS maintains various counters for ETH-SLM PDUs, which can be retrieved at any time for sessions that are initiated by a certain MEP. You can clear all the ETH-SLM statistics and PDU counters.

- **Support for Network Address Translation (ACX Series)**—Network Address Translation (NAT) is a method for modifying or translating network address information in packet headers. Either or both source and destination addresses in a packet may be translated. NAT can include the translation of port numbers as well as IP addresses. ACX Series routers support only source NAT for IPv4 packets. Static and destination NAT types are currently not supported on the ACX Series routers.



NOTE: In ACX Series routers, NAT is supported only on the ACX1100 AC-powered router.

- **Support for inline service interface (ACX Series)**—Junos OS for ACX Series Universal Access Routers support inline service interface. An inline service interface is a virtual physical interface that resides on the Packet Forwarding Engine. The **si-** interface makes it possible to provide NAT services without a special services PIC.

To configure inline NAT, you define the service interface as type **si-** (service-inline) interface. You must also reserve adequate bandwidth for the inline interface. This enables you to configure both interface or next-hop service sets used for NAT.



NOTE: In ACX Series routers, you can configure only one inline services physical interface as an anchor interface for NAT sessions: si-0/0/0.

- **Support for IPsec (ACX Series)**—You can configure IPsec on ACX Series Universal Access Routers. The IPsec architecture provides a security suite for the IP version 4 (IPv4) network layer. The suite provides functionality such as authentication of origin, data integrity, confidentiality, replay protection, and nonrepudiation of source. In addition to IPsec, Junos OS also supports the Internet Key Exchange (IKE), which defines mechanisms for key generation and exchange, and manages security associations. IPsec also defines a security association and key management framework that can be used with any network layer protocol. The security association specifies what protection policy to apply to traffic between two IP-layer entities. IPsec provides secure tunnels between two peers.



NOTE: IPsec is supported only on the ACX1100 AC-powered router.

- **Support for ATM OAM F4 and F5 cells (ACX Series)**—ACX Series routers provide Asynchronous Transfer Mode (ATM) support for the following Operations, Administration, and Maintenance (OAM) fault management cell types:
 - F4 alarm indication signal (AIS) (end-to-end)
 - F4 remote defect indication (RDI) (end-to-end)
 - F4 loopback (end-to-end)
 - F5 AIS
 - F5 RDI
 - F5 loopback

ATM OAM is supported on ACX1000, ACX2000, and ACX2100 routers, and on 16-port Channelized E1/T1 Circuit Emulation MICs on ACX4000 routers.

Junos OS supports the following methods of processing OAM cells that traverse through pseudowires with circuit cross-connect (CCC) encapsulation:

- Virtual path (VP) pseudowires (CCC encapsulation)
- Port pseudowires (CCC encapsulation)
- Virtual circuit (VC) pseudowires (CCC encapsulation)

For ATM pseudowires, the F4 flow cell is used to manage the VP level. On ACX Series routers with ATM pseudowires (CCC encapsulation), you can configure OAM F4 cell flows to identify and report virtual path connection (VPC) defects and failures. Junos OS supports three types of OAM F4 cells in end-to-end F4 flows:

- Virtual path AIS
- Virtual path RDI
- Virtual path loopback

For OAM F4 and F5 cells, IP termination is not supported. Also, Junos OS does not support segment F4 flows, VPC continuity check, or VP performance management functions.

For OAM F4 cells, on each VP, you can configure an interval during which to transmit loopback cells by including the **oam-period** statement at the **[edit interfaces interface-name atm-options vpi vpi-identifier]** hierarchy level. To modify OAM liveness values on a VP, include the **oam-liveness** statement at the **[edit interfaces interface-name atm-options vpi vpi-identifier]** hierarchy level.

- **Support for CESoPSN on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (ACX Series)**—You can configure structure-aware TDM CESoPSN on the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (model number: ACX-MIC-4COC3-1COC12CE) on ACX Series routers. This rate-selectable MIC can be configured as four OC3/STM1 ports or one OC12/STM4 port.
- **Support for Point-to-Point Protocol encapsulation (ACX Series)**—You can configure Point-to-Point Protocol (PPP) encapsulation on physical interfaces on ACX Series routers. PPP provides a standard method for transporting multiprotocol datagrams over a point-to-point link. PPP uses the High-Speed Data Link Control (HDLC) protocol for its physical interface and provides a packet-oriented interface for the network-layer protocols.

PPP is supported on the following MICs on ACX Series routers:

- On ACX1000 routers with 8-port built-in T1/E1 TDM MICs.
- On ACX2000 and ACX2100 routers with 16-port built-in T1/E1 TDM MICs.
- On ACX4000 routers with 16-port Channelized E1/T1 Circuit Emulation MICs.

On ACX Series routers, E1, T1, and NxDS0 interfaces support PPP encapsulation.

- **Support for Ethernet link aggregation (ACX Series)**—Junos OS for ACX Series Universal Access Routers support Ethernet link aggregation for Layer 2 bridging. Ethernet link aggregation is a mechanism for increasing the bandwidth of Ethernet links linearly and improving the links' resiliency by bundling or combining multiple full-duplex, same-speed, point-to-point Ethernet links into a single virtual link. The virtual link interface is referred to as a link aggregation group (LAG) or an aggregated Ethernet interface. The LAG balances traffic across the member links within an aggregated Ethernet interface and effectively increases the uplink bandwidth. Another advantage of link aggregation is increased availability, because the LAG is composed of multiple member links. If one member link fails, the LAG continues to carry traffic over the remaining links.
- **16-port Channelized E1/T1 Circuit Emulation MIC (ACX4000)**—ACX4000 Universal Access Routers support the 16-port Channelized E1/T1 Circuit Emulation MIC (model number ACX-MIC-16CHE1-T1-CE).

The key features supported on this MIC are:

- Structure-Agnostic TDM over Packet (SAToP)
- ATM encapsulation—Only the following ATM encapsulations are supported on this MIC:

- ATM CCC cell relay
- ATM CCC VC multiplex
- ATM pseudowires
- ATM quality-of-service (QoS) features—traffic shaping, scheduling, and policing
- ATM Operation, Administration, and Maintenance
- ATM (IMA) protocol at the T1/E1 level with up to 16 IMA (Inverse Multiplexing for ATM) groups. Each group can have 1-8 IMA links.
- **Support for PIM and IGMP in global domain (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports Protocol Independent Multicast (PIM) and Internet Group Management Protocol (IGMP) messages for multicast data delivery. ACX Series routers are used as a leaf in the multicast distribution tree so that subscribers in the global domain can directly connect to the ACX Series routers through IPv4 interfaces. ACX Series routers can also be used as a branch point in the tree so that they are connected to other downstream ACX Series or MX Series routers and send multicast data according to the membership established through the PIM or IGMP messaging.



NOTE: ACX Series routers support only sparse mode. Dense mode on ACX series is supported only for control multicast groups for autodiscovery of rendezvous point (auto-RP).

You can configure IGMP on the subscriber-facing interfaces to receive IGMP control packets from subscribers, which in turn triggers the PIM messages to be sent out of the network-facing interface toward the rendezvous point (RP).



NOTE: ACX Series routers do not support IPv6 interfaces for multicast data delivery and RP functionality.

- **Support for dying-gasp PDU generation (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports the generation of dying-gasp protocol data units (PDUs). Dying gasp refers to an unrecoverable condition such as a power failure. In this condition, the local peer informs the remote peer about the failure state. When the remote peer receives a dying-gasp PDU, it takes an action corresponding to the action profile configured with the **link-adjacency-loss** event.

ACX Series routers can generate and receive dying-gasp packets. When LFM is configured on an interface, a dying-gasp PDU is generated for the interface on the following failure conditions:
 - Power failure
 - Packet Forwarding Engine panic or a crash
- **Support for logical tunnels (ACX Series)**—Logical tunnel (**lt-**) interfaces provide quite different services depending on the host router. On ACX Series routers, logical tunnel interfaces enable you to connect a bridge domain and a pseudowire.

To create tunnel interfaces, an FPC and the corresponding Packet Forwarding Engine on an ACX Series router must be configured to be used for tunneling services at the **[edit chassis]** hierarchy level. The amount of bandwidth reserved for tunnel services must also be configured.

To create logical tunnel interfaces and the bandwidth in gigabits per second to reserve for tunnel services, include the **tunnel-services bandwidth (1g | 10g)** statement at the **[edit chassis fpc slot-number pic number]** hierarchy level.

- **Support for PPP encapsulation on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (ACX Series)**—On ACX4000 routers, you can configure Point-to-Point Protocol (PPP) encapsulation on physical interfaces on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP. PPP provides a standard method for transporting multiprotocol datagrams over a point-to-point link. PPP uses the High-Speed Data Link Control (HDLC) protocol for its physical interfaces and provides a packet-oriented interface for the network-layer protocols.

On ACX Series routers, E1, T1, and NxDS0 interfaces support PPP encapsulation.

IP class of service (CoS) is not supported on PPP interfaces. All the traffic is sent to the best effort queue (queue 0) and CoS code points are not processed. Also, fixed classifiers are not supported. PPP is supported only for IPv4 networks.

- **Support for dual-rate SFP+ modules (ACX Series)**—ACX2000, ACX2100, and ACX4000 routers support the dual-rate SFP+ optic modules. These modules operate at either 1 Gbps or 10 Gbps speeds. When you plug in the module to the small form-factor pluggable plus (SFP+) slot, the module can be set at either 1 Gbps or 10 Gbps.

ACX Series routers use the 2-port 10-Gigabit Ethernet (LAN) SFP+ MIC in the following two combinations:

- 2-port 10-Gigabit Ethernet (LAN) SFP+ uses BCM84728 PHY on ACX 2100/ACX4000 routers.
- 2-port 10-Gigabit Ethernet (LAN) SFP+ uses BCM8728/8747 on ACX2000 routers.

To configure an **xe** port in 1-Gigabit Ethernet mode, use the **set interfaces xe-x/y/z speed 1g** statement. To configure an **xe** port in 10-Gigabit Ethernet mode, use the **set interfaces xe-x/y/z speed 10g** statement. The default speed mode is 1-Gigabit Ethernet mode.

- **Support for inverse multiplexing for ATM (IMA) on Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (ACX Series)**—You can configure inverse multiplexing for ATM (IMA) on the Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC with SFP (model number: ACX-MIC-4COC3-1COC12CE) on ACX Series routers. You can configure four OC3/STM1 ports or one OC12/STM4 port on this rate-selectable MIC.
- **Support for TDR for diagnosing cable faults (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports Time Domain Reflectometry (TDR), which is a technology used for diagnosing copper cable states. This technique can be used to determine whether cabling is at fault when you cannot establish a link. TDR detects the defects by sending a signal through a cable, and reflecting it from the end of the

cable. Open circuits, short circuits, sharp bends, and other defects in the cable reflects the signal back at different amplitudes, depending on the severity of the defect. TDR diagnostics is supported only on copper interfaces and not on fiber interfaces.

TDR provides the following capabilities that you can use to effectively identify and correct cable problems:

- Display detailed information about the status of a twisted-pair cable, such as cable pair being open or short-circuited.
- Determine the distance in meters at which open or short-circuit is detected.
- Detect whether or not the twisted pairs are swapped.
- Identify the polarity status of the twisted pair.
- Determine any downshift in the connection speed.

Installation

- **Support for USB autoinstallation from XML file (ACX Series routers)**—Junos OS for ACX Series Universal Access Routers support USB autoinstallation using the configuration file in XML format. The USB-based autoinstallation process overrides the network-based autoinstallation process. If the ACX Series router detects a USB Disk-on-Key device containing a valid configuration file during autoinstallation, the router using the configuration file on Disk-on-Key instead of fetching the configuration from the network.
- **Support for hybrid mode of autoinstallation**—Junos OS for ACX Series Universal Access Routers support hybrid mode of autoinstallation. The autoinstallation mechanism allows the router to configure itself out-of-the-box with no manual intervention, using the configuration available on the network, locally through a removable media, or using a combination of both. ACX Series routers support the retrieval of partial configuration from an external USB storage device plugged into the router's USB port during the autoinstallation process. This partial configuration in turn facilitates the network mode of autoinstallation to retrieve the complete configuration file from the network. This method is called hybrid mode of autoinstallation.

Layer 2 Features

- **Support for Layer 2 security (ACX Series)**—ACX Series routers support bridge family firewall filters. These family filters can be configured at the logical interface level and can be scaled up to 124 terms for ingress traffic, and 126 terms for egress traffic.
- **Support for Ethernet Local Management Interface protocol (ACX Series)**—The Ethernet Local Management Interface (E-LMI) protocol on ACX Series Universal Access Routers supports Layer 2 circuit and Layer 2 VPN Ethernet virtual connection (EVC) types.

Junos OS for ACX Series Universal Access Routers support E-LMI only on provider edge (PE) routers.

- **Support for Layer 2 control protocols and Layer 2 protocol tunneling (ACX Series)**—You can configure spanning tree protocols to prevent Layer 2 loops in a bridge

domain. Layer 2 control protocols for ACX Series Universal Access Routers include the Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), VLAN Spanning Tree Protocol (VSTP), and Link Layer Discovery Protocol (LLDP). ACX Series routers can support up to 128 STP instances, which includes all instances of VSTP, MSTP, RSTP and STP.

Layer 2 protocol tunneling (L2PT) is supported on ACX Series routers. L2PT allows Layer 2 protocol data units (PDUs) to be tunneled through a network. L2PT can be configured on a port on a customer-edge router by using MAC rewrite configuration. MAC rewrite is supported for STP, Cisco Discovery Protocol (CDP), VLAN Trunk Protocol (VTP), IEEE 802.1X, IEEE 802.3ah, Ethernet Local Management Interface (E-LMI), Link Aggregation Control Protocol (LACP), Link Layer Discovery Protocol (LLDP), Multiple MAC Registration Protocol (MMRP), and Multiple VLAN Registration Protocol (MVRP) packets.

- **Support for Layer 2 bridging (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports Layer 2 bridging and Q-in-Q tunneling. A bridge domain is created by adding a set of Layer 2 logical interfaces in a bridge domain to represent a broadcast domain. Layer 2 logical interfaces are created by defining one or more logical units on a physical interface with encapsulation as **ethernet-bridge** or **vlan-bridge**. All the member ports of the bridge domain participate in Layer 2 learning and forwarding. You can configure one or more bridge domains to perform Layer 2 bridging. You can optionally disable learning on a bridge domain.



NOTE: ACX Series routers do not support the creation of bridge domains by using access and trunk ports.

On ACX Series routers, you can configure E-LAN and E-LINE services on bridge domains. When you configure E-LAN and E-LINE services by using a bridge domain without a **vlan-id** statement, the bridge domain should explicitly be normalized by an input VLAN map to a service VLAN ID and TPID. Explicit normalization is required when a logical interface's outer VLAN ID and TPID are not the same as the service VLAN ID and TPID of the service being configured.

- **Support for IEEE 802.1ad classifier (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports the IEEE 802.1ad classifier. Rewrite rules at the physical interface level support the IEEE 802.1ad bit value. The IEEE 802.1ad classifier uses IEEE 802.1p and DEI bits together. On logical interfaces, only fixed classifiers are supported.

You can configure either IEEE 802.1p or IEEE 802.1ad classifiers at the physical interface level. You can define the following features:

- IEEE 802.1ad classifiers (inner or outer)
- IEEE 802.1ad rewrites (outer)



NOTE: You cannot configure both IEEE 802.1p and IEEE 802.1ad classifiers together at the physical interface level.

ACX Series routers support the IEEE 802.1ad classifier and rewrite along with the existing class-of-service features for Layer 2 interfaces.

- **Support for OAM with Layer 2 bridging as a transport mechanism (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports the following OAM features that use Layer 2 bridging as a transport mechanism:
 - IEEE 802.3ah LFM—IEEE 802.3ah link fault management (LFM) operates at the physical interface level and the packets are sent using Layer 2 bridging as a transport mechanism.
 - Dying-gasp packets—Dying-gasp PDU generation operates at the physical interface level. Dying-gasp packets are sent through the IEEE 802.3ah LFM-enabled interfaces.
 - IEEE 802.1ag and ITU-T Y.1731 protocols on down MEPs—IEEE 802.1ag configuration fault management (CFM) and ITU-T Y.1731 performance-monitoring OAM protocols, which are used for end-to-end Ethernet services, are supported only on down maintenance association end points (MEPs). The ITU-T Y.1731 protocol supports delay measurement on down MEPs but does not support loss measurement on down MEPs.
- **Support for Storm Control**—Storm control is supported on ACX Series routers. Storm control is only applicable at the IFD level for ACX Series. When a traffic storm is seen on the interface configured for storm control, the default action is to drop the packets exceeding the configured bandwidth. No event is generated as part of this. Storm control is not enabled on the interface by default.
- **Support for RFC 2544-based benchmarking tests (ACX Series)**—Junos OS for ACX Series Universal Access Routers support RFC 2544-based benchmarking tests for E-LINE and ELAN services configured using bridge domains. RFC 2544 defines a series of tests that can be used to describe the performance characteristics of network interconnecting devices. RFC 2544 tests methodology can be applied to a single device under test, or a network service (set of devices working together to provide end-to-end service). When applied to a service, the RFC 2544 test results can characterize the service-level-agreement parameters.

RFC 2544 tests are performed by transmitting test packets from a device that functions as the generator or the initiator. These packets are sent to a device that functions as the reflector, which receives and returns the packets back to the initiator.

ACX Series routers support RFC 2544 tests to measure throughput, latency, frame loss rate, and back-to-back frames.

With embedded RFC 2544, an ACX Series router can be configured as an initiator and reflector.

- You can configure RFC 2544 tests on the following underlying services:
 - Between two IPv4 endpoints.

- Between two user-to-network interfaces (UNIs) of Ethernet Virtual Connection (EVC), Ethernet Private Line (EPL, also called E-LINE), Ethernet Virtual Private Line (EVPL), EVC (EPL, EVPL).
- **Support for IEEE 802.1ag and ITU-T Y.1731 OAM protocols on up MEPs (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports IEEE 802.1ag configuration fault management (CFM) and ITU-T Y.1731 performance-monitoring OAM protocols on up maintenance association end points (MEPs). CFM OAM protocol is supported on link aggregation group (LAG) or aggregated Ethernet (AE) interfaces. The ITU-T Y.1731 protocol supports delay measurement on up MEPs but does not support loss measurement on up MEPs.



NOTE: ACX Series routers do not support ITU-T Y.1731 OAM protocol on AE interfaces.

- **Support for Ethernet alarm indication signal (ACX Series)**—Junos OS for ACX Series Universal Access Routers support ITU-T Y.1731 Ethernet alarm indication signal function (ETH-AIS) to provide fault management for service providers. ETH-AIS enables you to suppress alarms when a fault condition is detected. Using ETH-AIS, an administrator can differentiate between faults at the customer level and faults at the provider level. When a fault condition is detected, a maintenance end point (MEP) generates ETH-AIS packets to the configured client levels for a specified duration until the fault condition is cleared. Any MEP configured to generate ETH-AIS packets signals to a level higher than its own. A MEP receiving ETH-AIS recognizes that the fault is at a lower level and then suppresses alarms at current level the MEP is in.

ACX Series routers support ETH-AIS PDU generation for server MEPs on the basis of the following defect conditions:

- Loss of connectivity (physical link loss detection)
- Layer 2 circuit or Layer 2 VPN down
- **Support for Ethernet ring protection switching (ACX Series)**--You can configure Ethernet ring protection switching (ERPS) on ACX Series routers to achieve high reliability and network stability. The basic idea of an Ethernet ring is to use one specific link, called the ring protection link (RPL), to protect the whole ring. Links in the ring will never form loops that fatally affect the network operation and services availability.

ACX Series routers support multiple Ethernet ring instances that share the physical ring. Each instance has its own control channel and a specific data channel. Each ring instance can take a different path to achieve load balancing in the physical ring. When no data channel is specified, ERP operates only on the VLAN ID associated with the control channel. G.8032 open rings are supported.

ACX Series routers do not support aggregate Ethernet-based rings.

To configure Ethernet ring protection switching, include the **protection-ring** statement at the **[edit protocols]** hierarchy level.

- **Support for integrated routing and bridging (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports integrated routing and bridging (IRB) functionality.

IRB provides routing capability on a bridge domain. To enable this functionality, you need to configure an IRB interface as a routing interface in a bridge domain and then configure a Layer 3 protocol such as IP or ISO on the IRB interface.

ACX Series routers support IRB for routing IPv4 packets. IPv6 and MPLS packets are not supported.

- **Support for IGMP snooping (ACX Series)**—Junos OS for ACX Series routers support IGMP snooping functionality. IGMP snooping functions by snooping at the IGMP packets received by the switch interfaces and building a multicast database similar to that a multicast router builds in a Layer 3 network. Using this database, the switch can forward multicast traffic only to the downstream interfaces of interested receivers. This technique allows more efficient use of network bandwidth, particularly for IPTV applications. You configure IGMP snooping for each bridge on the router.
- **Support for unicast reverse path forwarding (ACX Series)**—For interfaces that carry IPv4 or IPv6 traffic, you can reduce the impact of denial-of-service (DoS) attacks by configuring unicast reverse path forwarding (RPF). Unicast RPF helps determine the source of attacks and rejects packets from unexpected source addresses on interfaces where unicast RPF is enabled.

Reverse path forwarding is not supported on the interfaces that you configure as tunnel sources. This limitation affects only the transit packets exiting the tunnel.

To configure unicast reverse path forwarding, issue the **rpf-check** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet]** hierarchy level. RPF fail filters are not supported on ACX Series routers. The RPF check to be used when routing is asymmetrical is not supported.

- **Support for disabling local switching in bridge domains (ACX Series)**—In a bridge domain, when a frame is received from a customer edge (CE) interface, it is flooded to the other CE interfaces and all of the provider edge (PE) interfaces if the destination MAC address is not learned or if the frame is either broadcast or multicast.

To prevent CE devices from communicating directly include the **no-local-switching** statement at the **[edit bridge-domains *bridge-domain-name*]** hierarchy level. Configure the logical interfaces in the bridge domain as core-facing (PE interfaces) by including the **core-facing** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family*]** hierarchy level to specify that the VLAN is physically connected to a core-facing ISP router and ensure that the network does not improperly treat the interface as a client interface. When local switching is disabled, traffic from one CE interface is not forwarded to another CE interface.

- **Support for hierarchical VPLS (ACX Series)**—Hierarchical LDP-based VPLS requires a full mesh of tunnel LSPs between all the PE routers that participate in the VPLS service. Using hierarchical connectivity reduces signaling and replication overhead to facilitate large-scale deployments. In a typical IPTV solution, IPTV sources are in the public domain and the subscribers are in the private VPN domain.

For an efficient delivery of multicast data from the IPTV source to the set-top boxes or to subscribers in the private domain using the access devices (ACX Series routers in this case), P2MP LSPs and MVPN are necessary. Because VPLS and MVPN are not supported on ACX routers, an alternative approach is used to achieve hierarchical VPLS

(HPVLS) capabilities. The subscriber devices are connected to a VPLS or a Layer 3 VPN domain on the ACX Series (access) router and they are configured to import the multicast routes. The support for PIM snooping in Layer 3 interfaces, IGMP snooping in Layer 2 networks, IRB interfaces, and logical tunnel interfaces enables HPVLS support.

Management

- **Support for real-time performance monitoring (ACX Series)**—Real-time performance monitoring (RPM) allows you to perform service-level monitoring. When RPM is configured on a router, the router calculates network performance based on packet response time, jitter, and packet loss. You can configure these values to be gathered by HTTP, Internet Control Message Protocol (ICMP), TCP, and UDP requests. The router gathers RPM statistics by sending out probes to a specified probe target, identified by an IP address. When the target receives a probe, it generates responses that are received by the router. You set the probe options in the **test test-name** statement at the **[edit services rpm probe owner]** hierarchy level. You use the **show services rpm probe-results** command to view the results of the most recent RPM probes.



NOTE: Packet Forwarding Engine timestamping is available only for ICMP probes and for UDP probes with the destination port set to UDP_ECHO port (7).

- **Support for Virtual Router Redundancy Protocol version 2 (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports Virtual Router Redundancy Protocol (VRRP) version 2 configuration. VRRP enables hosts on a LAN to make use of redundant routers on that LAN without requiring more than the static configuration of a single default route on the hosts. Routers running VRRP share the IP address corresponding to the default route configured on the hosts. At any time, one of the routers running VRRP is the master (active) and the others are backups. If the master fails, one of the backup routers becomes the new master router, providing a virtual default router and enabling traffic on the LAN to be routed without relying on a single router. Using VRRP, a backup router can take over a failed default router within a few seconds. This is done with minimum VRRP traffic and without any interaction with the hosts.
- **Support for DHCP client and DHCP server (ACX Series)**—ACX Series Universal Access Routers can be enabled to function as a DHCP client and an extended DHCP local server. An extended DHCP local server provides an IP address and other configuration information in response to a client request in the form of an address-lease offer. An ACX Series router configured as a DHCP client can obtain its TCP/IP settings and the IP address from a DHCP local server.
- **Support for preserving DHCP server subscriber information (ACX Series)**—Junos OS for ACX Series Universal Access Routers preserves DHCP server subscriber binding information. ACX series router functioning as a DHCP server stores the subscriber binding information to a file and when the router reboots, the subscriber information is read from the file and restored.
- **Support for Two-Way Active Measurement Protocol (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports Two-Way Active Measurement Protocol (TWAMP). TWAMP provides a method for measuring round-trip IP performance

between two devices in a network. ACX Series routers support only the reflector side of TWAMP.

Routing

- **Support for ECMP flow-based forwarding (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports equal-cost multipath (ECMP) flow-based forwarding. An ECMP set is formed when the routing table contains multiple next-hop addresses for the same destination with equal cost. If there is an ECMP set for the active route, Junos OS uses a hash algorithm to choose *one* of the next-hop addresses in the ECMP set to install in the forwarding table. You can configure Junos OS so that multiple next-hop entries in an ECMP set are installed in the forwarding table. On ACX Series routers, per-flow load balancing can be performed to spread traffic across multiple paths between the routers.

ECMP flow-based forwarding is supported for IPv4, IPv6, and MPLS packets.

Security

- **Support for IP and MAC address validation (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports IP and MAC address validation. This feature enables the ACX Series router to validate that received packets contain a trusted IP source and an Ethernet MAC source address. Configuring MAC address validation can provide additional validation when subscribers access billable services. MAC address validation provides additional security by enabling the router to drop packets that do not match, such as packets with spoofed addresses.
- **Support for unattended boot mode (ACX Series)**—Junos OS for ACX Series Universal Access Routers support unattended boot mode. Unattended boot mode feature blocks any known methods to get access to the router from CPU reset till Junos OS login prompt, thereby preventing a user to make any unauthorized changes on the router such as viewing, modifying, or deleting configuration information.

Subscriber Access Management

- **Support for DHCP relay agent (ACX Series)**—You can configure extended DHCP relay options on an ACX Series router and enable the router to function as a DHCP relay agent. A DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server that might or might not reside in the same IP subnet.

To configure the DHCP relay agent on the router for IPv4 packets, include the **dhcp-relay** statement at the **[edit forwarding-options]** hierarchy level. You can also include the **dhcp-relay** statement at the **[edit routing-instances routing-instance-name forwarding-options]** and the **[edit routing-instances routing-instance-name protocols vrf]** hierarchy levels.

Timing and Synchronization

- **Support for PTP over Ethernet (ACX Series)**—Precision Time Protocol (PTP) is supported over IEEE 802.3 or Ethernet links on ACX Series routers. This functionality is supported in compliance with the IEEE 1588-2008 specification. PTP over Ethernet

enables effective implementation of packet-based technology that enables the operator to deliver synchronization services on packet-based mobile backhaul networks that are configured in Ethernet rings. Deployment of PTP at every hop in an Ethernet ring using the Ethernet encapsulation method enables robust, redundant, and high-performance topologies to be created that enables a highly-precise time and phase synchronization to be obtained.

- **PTP slave performance metrics (ACX Series)**—Precision Time Protocol (PTP) slave devices are used to provide frequency and time distribution throughout large networks. On ACX Series routers, PTP slave devices calculate performance metrics based on standard PTP timing messages. These performance metrics include both inbound and outbound packet delay and jitter between the PTP slave and master. Metrics are exported every 15 minutes to Junos Space. Performance metrics are also stored locally on the ACX Series router and can be accessed with the **show ptp performance-monitor [short-term | long-term]** command.
- **Support for hybrid mode (ACX Series)**—Junos OS for ACX Series Universal Access Routers supports hybrid mode, which is a combined operation of Synchronous Ethernet and Precision Time Protocol (PTP). In hybrid mode, the synchronous Ethernet equipment clock (EEC) on the router derives the frequency from Synchronous Ethernet and the phase and time of day from PTP. Time synchronization includes both phase synchronization and frequency synchronization.

Synchronous Ethernet supports hop-by-hop frequency transfer, where all interfaces on the trail must support Synchronous Ethernet. PTP (also known as IEEE 1588v2) synchronizes clocks between nodes in a network, thereby enabling the distribution of an accurate clock over a packet-switched network.

To configure the router in hybrid mode, you must configure Synchronous Ethernet options at the **[edit chassis synchronization]** hierarchy level and configure PTP options at the **[edit protocols ptp]** hierarchy level. Configure hybrid mode options by including the **hybrid** statement at the **[edit protocols ptp slave]** hierarchy level.

Related Documentation

- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

Changes in Default 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 15.1R3 for the ACX Series Universal Access Routers.

Interfaces and Chassis

- **Connectivity fault management MEPs on Layer 2 circuits and Layer 2 VPNs**—On interfaces configured on ACX Series routers, you no longer need to configure the **no-control-word** statement at either the **[edit protocols l2circuit neighbor *neighbor-id* interface *interface-name*]** or the **[edit routing-instances *routing-instance-name* protocols l2vpn]** hierarchy level for Layer 2 circuits and Layer 2 VPNs over which you are running CFM maintenance association end points (MEPs). This configuration is not needed because ACX Series routers support the control word for CFM MEPs. The control word is enabled by default.
- In the output of the **show interfaces** command under the **MAC Statistics** section, any packet whose size exceeds the configured MTU size is considered as an oversized frame and the value displayed in the **Oversized frames** field is incremented. The value displayed in the **Jabber frames** field is incremented when a bad CRC frame size is between 1518 bytes and the configured MTU size.
- **Support for chained composite next hop in Layer 3 VPNs**—Next-hop chaining (also known as chained composite next hop) is a composition function that concatenates the partial rewrite strings associated with individual next hops to form a larger rewrite string that is added to a packet. To configure the router to accept up to one million Layer 3 VPN route updates with unique inner VPN labels, include the **l3vpn** statement at the **[edit routing-options forwarding-table chained-composite-next-hop ingress]** hierarchy level. The **l3vpn** statement is disabled by default.

Related Documentation

- [New and Changed Features on page 7](#)
- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Documentation Updates on page 34](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

Known Behavior

There are no known limitations in Junos OS Release 15.1R3 for the ACX Series.

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

Related Documentation

- [New and Changed Features on page 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Documentation Updates on page 34](#)

- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R3 for the ACX Series Universal Access Routers.

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 on page 25](#)
- [Firewall Filters on page 26](#)
- [Interfaces and Chassis on page 28](#)
- [Integrated Routing and Bridging on page 30](#)
- [Layer 2 Services on page 31](#)
- [MPLS Applications on page 32](#)
- [Network Management on page 32](#)
- [Statistics on page 33](#)
- [Timing and Synchronization on page 33](#)

Class of Service

- When the **rewrite-rules** statement is configured with the **dscp** or the **inet-precedence** options at the [**edit class-of-service interfaces**] hierarchy level, the expectation is that the DiffServ code point (DSCP) or IPv4 precedence rewrite rules take effect only on IP packets. However, in addition to the IP packets, the DSCP or IPv4 rewrite takes effect on the IP header inside the Ethernet pseudowire payload as well. This is not applicable for ACX4000 router. [PR664062](#)
- In an ACX4000 router, whenever the scheduling and shaping parameters of a port or any of its queues are changed, the entire scheduling configuration on the port is erased and the new configuration is applied. During the time when such a configuration change is taking place, the traffic pattern does not adhere to user parameters. It is recommended that the scheduling configurations are done much earlier before live traffic. [PR840313](#)
- The VLAN packet loss priority (PLP) is incorrectly set when untagged VLAN frames are received on the ingress interface with DSCP or IP precedence classification enabled and the NNI (egress) interface does not contain IEEE 802.1p rewrite rules. [PR949524](#)
- On the ACX4000 router, when class of service is not configured, traffic egressing out of the UNI port is going through all the queues instead of a default queue with code point 000. This issue is seen with the 500 pseudowire. As a workaround, you can use the following CLI command to avoid this issue:

```
user@host# set class-of-service system-defaults classifiers exp default
PR1123122
```

CoS limitations on PPP and MLPPP interfaces

The following are the common limitations on PPP and MLPPP interfaces:

- Traffic loss is observed when a CoS configuration is changed.
- Scheduling and shaping feature is based on CIR-EIR model and not based on weighted fair queuing (WFQ) model.
- The minimum transmit rate is 32 Kbps and the minimum supported rate difference between transmit rate and shaping rate is 32 Kbps.
- Buffer size is calculated based on the average packet size of 256 bytes.
- **Low** and **High** are the only loss priority levels supported.
- The mapping between forwarding class and queue is fixed as follows:
 - **best-effort** is queue 0
 - **expedited-forwarding** is queue 1
 - **assured-forwarding** is queue 2
 - **network-control** is queue 3

The following are the specific CoS limitations on MLPPP interfaces:

- Percentage rate configuration is not supported for shaping and scheduling. Rate configuration is only supported in terms of bits per second.
- Buffer size is calculated based on a single member link (T1/E1) speed and is not based on the number of member links in a bundle.
- Supports only **transmit-rate exact** configuration without fragmentation-map. Shaping and priority will not be supported without fragmentation-map.
- If fragmentation-map configured, shaping is supported on forwarding class with different priorities. If two or more forwarding classes are configured with the same priority, then only **transmit-rate exact** is supported for the respective forwarding class.
- Supports only one-to-one mapping between a forwarding class and a multiclass. A forwarding class can only send traffic corresponding to one multiclass.

The following is the specific CoS limitation on PPP interfaces:

- The distribution of excess rate between two or more queues of same priority happens on a first-come first-served basis. The shaping rate configured on the respective queue remains valid.

Firewall Filters

- In ACX Series routers, the following Layer 2 control protocols packet are not matched (with **match-all** term) by using the bridge family firewall filter applied on a Layer 2 interface:

- Slow-Protocol/LACP MAC (01:80:c2:00:00:02)
- E-LMI MAC ((01:80:c2:00:00:07)
- IS-IS L2 MAC (01:80:c2:00:00:14/09:00:2B:00:00:14)
- STP BPDU (01:80:c2:00:00:00)
- VSTP BPDU (01:00:0C:CC:CC:CD)
- LLDP/PTP (01:80:c2:00:00:0E)

When layer rewrite is configured:

- VTP/CDP (01:00:0C:CC:CC:CC)
- L2PT RW MAC (01:00:0C:CD:CD:D0)
- MMRP (01:80:C2:00:00:20)
- MVRP (01:80:C2:00:00:21)

As a workaround, to match the Layer 2 control packet flows with a bridge family filter term, you must explicitly specify the destination MAC match (along with other MAC matches) in the firewall filter term and in the match term. [PR879105](#)

- In ACX Series routers, a firewall filter cannot be applied to a logical interface configured with **vlan-id-list** or **vlan-range**. As a workaround, you can configure the interface-specific statement, which can be applied to the **bridge**, **inet**, or **mpls** family firewall filter. [PR889182](#)
- In ACX Series routers, packet drops in the egress interface queue are also counted as *input packet rejects* under the **Filter statistics** section in the output of the **show interface input-interfaces extensive** command when the command is run on the ingress interface. [PR612441](#)
- When the **statistics** statement is configured on a logical interface—for example, **[edit interface name-X unit unit-Y]**; the **(policer | count | three-color-policer)** statements are configured in a firewall filter for the **family any**—for example, **[edit firewall family any filter filter-XYZ term term-T then]** hierarchy level; and the configured **filter-XYZ** is specified in the **output** statement of the logical interface at the **[edit interface name-X unit unit-Y filter]** hierarchy level, the counters from the configuration of another firewall family filter on the logical interface do not work. [PR678847](#)
- The policing rate can be incorrect if the following configurations are applied together:
 - The **policer** or **three-color-policer** statement configured in a firewall filter—for example, **filter-XYZ** at the **[edit firewall family any filter filter-XYZ term term-T then]** hierarchy level, and **filter-XYZ** is specified as an ingress or egress firewall filter on a logical interface—for example, **interface-X unit-Y** at the **[edit interface interface-X unit unit-Y filter (input|output) filter-XYZ]** hierarchy level.
 - The **policer** or **three-color-policer** statement configured in a firewall filter—for example, **filter-ABC** at the **[edit firewall family name-XX filter filter-ABC term term-T then]** hierarchy level, and **filter-ABC** is configured as an ingress or egress firewall filter on a family of the same logical interface **interface-X unit-Y** at the **[edit interface interface-X unit unit-Y family name-XX filter (input|output) filter-ABC]** hierarchy level.



NOTE: If one of these configurations is applied independently, then the correct policer rate can be observed.

PR678950

Interfaces and Chassis

- Egress maximum transmission unit (MTU) check value of an interface is different for tagged and untagged packets. If an interface is configured with CLI MTU value as x , then the following would be the checks depending on outgoing packet type:
 - Egress MTU value for untagged packet = $x - 4$
 - Egress MTU value for single-tagged packet = x
 - Egress MTU value for double-tagged packet = $x + 4$



NOTE: The ingress MTU check is the same for all incoming packet types.

There is no workaround available. [PR891770](#)

- In ACX Series routers, when STP is configured on an interface, the detailed interface traffic statistics show command output does not show statistics information but displays the message **Dropped traffic statistics due to STP State**. However, the drop counters are updated. There is no workaround available. [PR810936](#)
- When the **differential-delay number** option is configured in the **ima-group-option** statement at the **[edit interfaces at-fpc/pic/ima-group-no]** hierarchy level, with a value less than 10, some of the member links might not come up and the group might remain down resulting in traffic loss. A workaround is to keep the differential delay value above 10 for all IMA bundles. [PR726279](#)
- The ACX Series routers support logical interface statistics, but do not support the address family statistics. [PR725809](#)
- BERT error insertion and bit counters are not supported by the IDT82P2288 framer. [PR726894](#)
- All 4x supported TPIDs cannot be configured on different logical interfaces of a physical interface. Only one TPID can be configured on all logical interfaces of a physical interface. But different physical interfaces can have different TPIDs. As a workaround, use TPID rewrite. [PR738890](#)
- The ACX Series routers do not support logical interface statistics for logical interfaces with **vlan-list** or **vlan-range** configured. [PR810973](#)
- CFM up-MEP session (to monitor pseudowire service) does not come up when output VLAN map is configured as **push** on AC logical interface. This is due to a hardware limitation in the ACX4000 router. [PR832503](#)

- For ATM interfaces with **atm-ccc-cell-relay** and **atm-ccc-vc-mux** encapsulation types configured, and with shaping profile configured on the interfaces, traffic drop is observed when the configured shaping profile is changed. This problem occurs with 16-port Channelized E1/T1 Circuit Emulation MICs on ACX4000 routers. As a workaround, you must stop the traffic on the Layer 2 circuit before changing any of the traffic shaping profile parameters. [PR817335](#)
- In the case of normalized bridge domain, with double-tagged aggregated Ethernet interface as ingress, the classification based on inner tag does not work for ACX4000. To do classification based on inner tag, configure the bridge domain with explicit normalization and configure input and output VLAN map to match the behavior. [PR869715](#)
- The MAC counter behavior of 10-Gigabit Ethernet is different compared to 1-Gigabit Ethernet.

On 1-Gigabit Ethernet interfaces, if the packet size is greater than 1518 bytes, irrespective of whether the packet is tagged or untagged, the **Oversized** counter gets incremented. If the packet has a CRC error, then the **Jabber** counter gets incremented.

On 10-Gigabit Ethernet interfaces, if the packet size is greater than 1518 bytes and the packet is untagged, then the **Oversized** counter gets incremented. If the packet has a CRC error, then the **Jabber** counter gets incremented.

If the packet is tagged (TPID is 0x8100), then the **Oversized** counter is incremented only if the packet size is greater than 1522 bytes (1518 + 4 bytes for the tag). The **Jabber** counter is incremented only if the packet size is greater than 1522 bytes and the packet has a CRC error.

The packet is considered as tagged if the outer TPID is 0x8100. Packets with other TPIDs values (for example, 0x88a8, 0x9100, or 0x9200) are considered as untagged for the counter. There is no workaround available. [PR940569](#)

- Layer 2 RFC2544 benchmarking test cannot be configured to generate dual-tagged frames when the UNI interface is configured for the QnQ service. This occurs when the input VLAN map **push** is configured on the UNI interface. There is no workaround available. [PR946832](#)
- After running RFC2544 tests, PTP stops working when the tests are performed on the same router. A workaround is to reboot FEB after running the RFC2544 tests. [PR944200](#)
- When an ACX1100 router with AC power is configured as PTP slave or boundary clock, the router does not achieve PTP accuracy within the specification (1.5 us), even if the PTP achieves the state **Phase Aligned**. [PR942664](#)
- Layer 2 RFC2544 benchmark test fails for packet sizes 9104 and 9136 when the test bandwidth is less than 10-MB and the NNI interface link speed is 10-MB. This behavior is also seen when the 10-MB policer or shaper is configured on the NNI interface. The issue will not be seen if the egress queue is configured with sufficient queue buffers. [PR939622](#)
- **Limitations on logical tunnel interfaces**—The following limitations apply when you configure logical tunnel (LT) interfaces in ACX Series Universal Access Routers:

- ACX router supports a total of two LT interfaces in a system, one of bandwidth 1G and another of bandwidth 10G.
- The bandwidth configured on the LT interface is shared between upstream and downstream traffic on that interface. The effective available bandwidth for the service is half the configured bandwidth.
- Supported encapsulations on LT interface are **ethernet-bridge**, **ethernet-ccc**, **vlan-bridge**, **vlan-ccc**.
- Total number of LT logical interfaces supported on a router is 30.
- If an LT interface with bandwidth 1G is configured and port-mirroring is also configured on the router, then LT physical interface statistics may not be accurate for that LT interface.
- Default classifiers are not available on the LT interface if a non-Ethernet PIC is used to create the LT interface.
- LT interfaces do not support protocol configuration.

Integrated Routing and Bridging

The following are the limitations on integrated routing and bridging (IRB) for ACX Series Universal Access Routers.

At the IRB device level, the following limitations apply:

- Behavior aggregate (BA) classifiers are not supported
- Statistics are not supported.

On an IRB logical interface, the following limitations apply:

- Statistics and Layer 2 policers are not supported
- Only inet and iso families are supported

On an IRB logical interface family inet, the following limitations apply:

- Policers, rpf-check, and dhcp-client are not supported

When firewall is applied on an IRB logical interface family inet, the following limitations apply:

- Default (global) filters are not supported.
- Supports only accept, forwarding-class, and loss-priority actions.
- Supports only input filters

Interface Limitations—IRB configurations supports a maximum of 1000 logical interfaces on a box.

Class-of-service Limitations—The following are CoS limitations for IRB:

- Maximum of 16 fixed classifiers are supported. Each classifier consumes two filter entries and is shared with RFC 2544 sessions. Total number of shared filter entries is 32.
- Maximum of 64 multifield filter classifiers are supported. Each classifier takes two filter entries. Total 128 entries are shared between family inet based classifiers on IRB and normal Layer 3 logical interfaces.
- Maximum 24 forwarding class and loss priority combinations can be rewritten. Each rewrite rule takes single entry from egress filters. Total of 128 entries are shared by rewrite-rules and all other output firewall filters.
- IRB rewrite is supported only on the ACX4000 Series router.

Firewall Limitations—The following are the firewall limitations for IRB:

- IRB supports only family inet filters.
- Only interface-specific and physical-interface specific filters are supported.
- Only forwarding-class and loss-priority actions are supported, other actions are not supported.

Layer 2 Services

Limitations on Layer 2 bridging

The following Layer 2 bridging limitations apply for ACX Series Universal Access Routers:

- A bridge domain cannot have two or more logical interfaces that belong to the same physical interface.
- A bridge domain with dual VLAN ID tag is not supported.
- The following input VLAN map functions are not supported because the bridge domain should have a valid service VLAN ID after normalization:
 - **pop-pop** on double-tagged logical interface.
 - **pop** on a single-tagged logical interface.
 - VLAN map with VLAN ID value set to 0.
- **swap-push** and **pop-swap** VLAN map functions are not supported.
- The maximum number of supported input VLAN maps with TPID **swap** is 64.
- MAC learning cannot be disabled at the logical interface level.
- MAC limit per logical interface cannot be configured.
- All STP ports on a bridge domain must belong to the same MST (multiple spanning tree) instance.
- If a logical interface is configured with Ethernet bridge encapsulation with **push-push** as the input VLAN map, normalization does not work when single-tagged or double-tagged frames are received on the logical port. Untagged frames received on the logical interface are normalized and forwarded correctly.

- On a priority-tagged logical interface with the output VLAN map function **pop**, egress VLAN filter check does not work.
- Output VLAN map function **push** cannot work on a dual-tagged frame egressing a logical interface.
- In a bridge domain configured with **vlan-id** statement, when a dual-tagged frame enters a non-dual-tagged logical interface and exits a dual-tagged logical interface, the VLAN tags are not translated correctly at egress.

Limitations on integrated routing and bridging

The following integrated routing and bridging (IRB) limitations apply for ACX Series Universal Access Routers:

At the IRB device level, the following limitations apply:

- Behavior aggregate (BA) classifiers are not supported
- Statistics are not supported.

On an IRB logical interface, the following limitations apply:

- Statistics and Layer 2 policers are not supported
- Only inet and iso families are supported

On an IRB logical interface family inet, the following limitations apply:

- Policers, rpf-check, and dhcp-client are not supported

When firewall is applied on an IRB logical interface family inet, the following limitations apply:

- Default (global) filters are not supported.
- Supports only accept, forwarding-class, and loss-priority actions.
- Supports only input filters

MPLS Applications

- The scaling numbers for pseudowires and MPLS label routes published for the ACX Series routers are valid only when the protocols adopt graceful restart. In case of non-graceful restart, the scaling numbers would become half of the published numbers. [PR683581](#)

Network Management

- In a connectivity fault management (CFM) up-mep session, when a remote-mep error is detected, the local-mep does not set the RDI bit in the transmitted continuity check messages (CCM). This problem is not seen in ACX4000 routers and in down-mep sessions. There is no workaround available. [PR864247](#)

- The ACX Series routers do not support the configuration of RPM probes to a routing instance along with the configuration of the **hardware-timestamp** statement at the `[edit services rpm probe owner test test-name]` hierarchy level. [PR846379](#)

Statistics

- ACX Series routers do not support route statistics per next hop and per flow for unicast and multicast traffic. Only interface-level statistics are supported.
- The **show multicast statistics** command is not supported on ACX Series routers. [\[PR954273\]](#)

Timing and Synchronization

- When you use the **replace pattern** command to toggle from a secure slave to an automatic slave or vice versa in the PTP configuration of a boundary clock, the external slave goes into a freerun state. The workaround is to use the **delete** and **set** commands instead of the **replace pattern** command. [PR733276](#)
- When you configure PTP over IPv4 with a dual logical interface path on the same physical interface, some of the routers in the ring get stuck in a **FREERUN** mode. This happens while switching from a primary logical interface path to a secondary logical interface path. [PR1134121](#)

Related Documentation

- [New and Changed Features on page 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Behavior on page 24](#)
- [Resolved Issues on page 33](#)
- [Documentation Updates on page 34](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

Resolved Issues

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

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

- [Resolved Issues on page 33](#)

Resolved Issues

There are no resolved issues in 15.1R3.

Related Documentation

- [New and Changed Features on page 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)

- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Documentation Updates on page 34](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

Documentation Updates

There are no errata or changes in Junos OS Release 15.1R3 for the ACX documentation.

Related Documentation

- [New and Changed Features on page 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)
- [Product Compatibility on page 35](#)

Migration, Upgrade, and Downgrade Instructions

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Universal Access Routers. 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 34](#)

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release, even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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

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

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

**Related
Documentation**

- [New and Changed Features on page 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Documentation Updates on page 34](#)
- [Product Compatibility on page 35](#)

Product Compatibility

- [Hardware Compatibility on page 35](#)

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 ACX Series routers 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 7](#)
- [Changes in Default Behavior and Syntax on page 23](#)
- [Known Behavior on page 24](#)
- [Known Issues on page 25](#)
- [Resolved Issues on page 33](#)
- [Documentation Updates on page 34](#)
- [Migration, Upgrade, and Downgrade Instructions on page 34](#)

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 15.1R3 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at <http://www.juniper.net/techpubs/software/junos/>.

- [New and Changed Features on page 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R3 for the EX Series.



NOTE: The following EX Series platforms are supported in Junos OS Release 15.1R3: EX2200, EX3300, EX4200, EX4300, EX4500, EX4550, EX4600, EX6200, EX8200, and EX9200.



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

- The J-Web Platform package—Installed as part of Junos OS; provides basic functionalities of J-Web.
- The 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, EX4300, EX4500, EX4550, and EX6200 Junos OS Release 15.1R1 install images.

For details about the J-Web distribution model, see [Release Notes: J-Web Application Package Release 15.1A2 for Juniper Networks EX Series Ethernet Switches](#).

- [Hardware on page 37](#)
- [Authentication and Access Control on page 38](#)
- [Interfaces and Chassis on page 39](#)
- [Junos OS XML API and Scripting on page 40](#)
- [Management on page 40](#)
- [MPLS on page 41](#)
- [Port Security on page 41](#)
- [Software Installation and Upgrade on page 42](#)
- [Spanning-Tree Protocols on page 42](#)

Hardware

- **EX9200-MPC line card for EX9200 switches**—Starting with Junos OS Release 15.1R3, EX9200 switches support the new EX9200-MPC line card. It is a modular line card that has two slots on the faceplate in which you can install any of the following modular interface cards (MICs):
 - EX9200-10XS-MIC: It has ten 10-Gigabit Ethernet small form-factor pluggable plus (SFP+) ports, which can house SFP+ transceivers. These ports support 10GBASE-SR, 10GBASE-LR, 10GBASE-ER, and 10GBASE-ZR transceivers.
 - EX9200-20F-MIC: It has twenty 1-Gigabit Ethernet small form-factor pluggable (SFP) ports with Media Access Control Security (MACsec) capability, each of which can house 1-gigabit SFP transceivers. These ports support 1000BASE-T, 1000BASE-SX, 100BASE-FX, 1000BASE-LX, 1000BASE-BX-U, 1000BASE-BX-D, 100BASE-BX-U, 100BASE-BX-D, and 1000BASE-LH transceivers.
 - EX9200-40T-MIC: It has 40 RJ-45 ports.

You can install the MICs in the following configurations:

- One EX9200-10XS-MIC
- One EX9200-20F-MIC
- One EX9200-10XS-MIC and one EX9200-20F-MIC
- Two EX9200-10XS-MICs
- Two EX9200-20F-MICs
- One EX9200-40T-MIC

You can transmit up to 130 gigabits of traffic through the line card without packet drop.

- **New optical transceiver support**—Starting with Junos OS Release 15.1R3, the 40-Gigabit Ethernet quad small form-factor pluggable plus (QSFP+) ports on EX9200-4QS and EX9200-6QS line cards for EX9200 switches support the transceiver JNP-QSFP-40G-LX4.

Authentication and Access Control

- **Central Web authentication (EX2200, EX3300, and EX4300)**—Starting with Junos OS Release 15.1R3, you can configure central web authentication to redirect Web browser requests to a login page that requires the user to input a username and password. Upon successful authentication, the user is allowed to access the network. The login process is handled by a central web authentication server, which provides scaling benefits over local web authentication, also known as captive portal.

Central web authentication is useful for providing network access to temporary users, such as visitors to a corporate site, who are trying to access the network using devices that are not 802.1X-enabled. Web authentication can also be used as a fall back authentication method for regular network users who have 802.1X-enabled devices, but fail authentication because of other issues, such as expired network credentials.

See [Understanding Central Web Authentication](#).

- **RADIUS-initiated changes to an authorized user session (EX2200, EX3300, and EX4300)**—Starting with Junos OS Release 15.1R3, EX2200, EX3300, and EX4300 switches support changes to an authorized user session that are initiated by the authentication server. The server can send the switch a Disconnect message to terminate the session, or a Change of Authorization (CoA) message to modify the session authorization attributes. CoA messages are typically used to change data filters or VLANs for an authenticated host.

See [Understanding RADIUS-Initiated Changes to an Authorized User Session](#).

- **Flexible authentication order (EX2200, EX3300, and EX4300)**—Starting with Junos OS Release 15.1R3, you can configure the order of authentication methods that the switch will use to authenticate an end device. By default, the switch will first attempt to authenticate using 802.1X authentication, then MAC RADIUS authentication, and then captive portal. You can override the default order of authentication methods by configuring the **authentication-order** statement to specify that the switch use either

802.1X authentication or MAC RADIUS authentication first. Captive portal must always be last in the order of authentication methods.

See [Understanding Authentication on EX Series Switches](#).

- **RADIUS accounting interim updates (EX4300)**—Starting with Junos OS Release 15.1R3, you can configure and EX4300 switch to send periodic updates for a user accounting session at a specified interval to the accounting server. Interim accounting updates are included in the exchange of messages between the client and the accounting server. In RADIUS accounting, the client sends Accounting-Request messages to the server, which acknowledges receipt of the requests with Accounting-Response messages. Interim accounting updates are sent in Accounting-Request messages with the Acct-Status-Type set to Interim-Update.

See [Understanding 802.1X and RADIUS Accounting on EX Series Switches](#).

- **Support for multiple terms in a filter sent from the RADIUS server (EX4300)**—Starting with Junos OS Release 15.1R3, you can use RADIUS server attributes to implement dynamic firewall filters with multiple terms on a RADIUS authentication server. These filters can be dynamically applied on all switches that authenticate supplicants through that server, eliminating the need to configure the same filter on multiple switches. You can define the filters directly on the server by using the Juniper-Switching-Filter attribute, which is a RADIUS attribute specific to Juniper Networks, also known as a vendor-specific attribute (VSA). Filter terms are configured using one or more match conditions and a resulting action.

See [Understanding Dynamic Filters Based on RADIUS Attributes](#).

- **EAP-PAP protocol support for MAC RADIUS authentication (EX2200, EX3300, and EX4300)**—Starting with Junos OS Release 15.1R3, you can configure the switch to use the Password Authentication Protocol (PAP) when authenticating clients with the MAC RADIUS authentication method. PAP transmits plaintext passwords over the network without encryption. It is required for use with LDAP (Lightweight Directory Access Protocol), which supports plaintext passwords for client authentication. This feature is configured by using the **authentication-protocol** CLI statement at the **[edit protocols dot1x authenticator interface interface-name mac-radius]** hierarchy level.

See [Understanding Authentication on EX Series Switches](#).

Interfaces and Chassis

- **LACP minimum link support on LAGs (EX9200)**—Starting with Junos OS Release 15.1R3, LACP minimum link support is added to the existing minimum link feature. The minimum-link configuration specifies that a required minimum bandwidth is provided for LAG interfaces. When there are not enough active links to provide this minimum bandwidth for a LAG interface, the LAG interface is brought down. The LACP minimum-link feature enhances the existing minimum-link feature by bringing down the LAG interface on the peer device as well as on the device on which you have configured minimum links. Before the LACP minimum link enhancement was made, if you configured the minimum link feature on one device but could not or had not configured it on the peer device, traffic would exit the LAG interface on the peer device although it would be dropped at the destination because the LAG interface on the peer

is not be brought down. LACP minimum link is enabled by default when you configure minimum links.

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

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

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

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

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

Junos OS XML API and Scripting

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

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

Management

- **Support for YANG features, including configuration hierarchy must constraints published in YANG, and a module that defines Junos OS YANG extensions (EX**

Series)—Starting with Junos OS Release 15.1, the Juniper Networks **configuration** YANG module includes configuration constraints published using either the YANG **must** statement or the Junos OS YANG extension **junos:must**. Constraints that cannot be mapped directly to the YANG **must** statement, which include expressions containing special keywords or symbols such as **all**, **any**, **unique**, **\$**, **_**, and wildcard characters, are published using **junos:must**.

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

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

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

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

MPLS

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

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

Port Security

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

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

- **MAC move limiting support (EX9200 switches)**—Starting with Junos OS Release 15.1R1, MAC move limiting is supported on EX9200 switches. MAC move limiting provides port security by controlling the number of MAC address moves that are allowed in a

VLAN in one second. When MAC move limiting is configured, the switch tracks MAC address movements on access and trunk interfaces. A MAC address move occurs when an interface on the switch receives a packet with a source MAC address that has already been learned by the switch, but on a different interface. If a MAC address moves more than the configured number of times within one second, you can configure an action to be taken on incoming packets with new source MAC addresses. The incoming packets can be dropped, logged or ignored. You can also specify an action to shutdown or temporarily disable the interfaces associated with that MAC address.

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

Software Installation and Upgrade

- **Support for FreeBSD 10 kernel for Junos OS (EX9200 switches)**—Starting with Junos OS Release 15.1, on EX9200, FreeBSD 10 is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. Because installation restructures the file system, logs and configurations are lost unless precautions are taken. There are now Junos OS and OAM volumes, which provide the ability to boot from the OAM volume upon failures. Some system commands display a different output than on earlier releases and a few others are deprecated.

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

Spanning-Tree Protocols

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

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

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

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

Related Documentation

- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)

- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 15.1R3 for the EX Series.

- [Dynamic Host Configuration Protocol on page 43](#)

Dynamic Host Configuration Protocol

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

Related Documentation

- [New and Changed Features on page 36](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R3 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](#)
- [Interfaces and Chassis](#)
- [J-Web](#)
- [Network Management and Monitoring](#)
- [Port Security](#)
- [Software Installation and Upgrade](#)
- [Spanning-Tree Protocols](#)
- [Virtual Chassis](#)

Authentication and Access Control

- On EX9200 switches, if you configure a firewall filter such that the number of characters in the filter name, term name, and counter name added together exceeds 128 characters,

802.1X (dot1x) authentication might fail and cause the Network Processing Card (NPC) to crash. As a workaround, configure the filter name, term name, and counter name such that when the sum of the number of characters in those three names is added to the sum of the number of characters in the interface name and the MAC address, the total does not exceed 128. [PR1083132](#)

- On EX9200 switches, 802.1X (dot1x) authentication might not be performed if a voice VLAN is changed or modified to a data VLAN after a client is authenticated in that voice VLAN. This problem occurs when a VoIP VLAN is configured, a client is authenticated in a configured data VLAN, and then the VoIP VLAN is configured as a new data VLAN (that is, you delete the VoIP configuration and delete the current data VLAN membership, and configure the original VoIP VLAN as the new data VLAN). [PR1074668](#)

Interfaces and Chassis

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

J-Web

- 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](#)
- If you access the J-Web interface using Microsoft Internet Explorer version 7, on the BGP Configuration page (Configure > Routing > BGP), all flags might be shown in the Configured Flags list (in the Edit Global Settings window, in the Trace Options tab), even though the flags are not configured. As a workaround, use the Mozilla Firefox browser. [PR603669](#)
- On the J-Web interface, on the Route Information page (Monitor > Routing > Route Information), the Next Hop column displays only the interface address, and the corresponding IP address is missing. The title of the first column displays **Static Route Address** instead of **Destination Address**. As a workaround, use the **show route detail** CLI command to fetch the IP address of the next-hop interface. [PR684552](#)
- On the J-Web interface, HTTPS access might work even 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 be displayed as up in the J-Web dashboard. [PR862411](#)

Network Management and Monitoring

- On EX4300 switches, if you configure a remote analyzer with an output IP address that is reachable through routes learned by BGP, the analyzer state is DOWN. [PR1007963](#)

Port Security

- On EX9200 switches, a DHCPv6 security dynamic entry binding might not work properly on an IPv6 IRB interface that is linked to a DHCP snooping VLAN. [PR1059623](#)
- On EX2200 switches, if you issue the **request system services dhcp release *interface-name*** operational command, an IP address release message DHCP packet is sent from the client and processed at the server. When the client clears the IP address on the same interface, the kernel generates an event message, which is processed at the client and triggers the DHCP client state machine, which leads to the interface acquiring a new IP address from the server. If you then issue the **show system services dhcp client *interface-name*** command, the output of that command indicates that the issued **request system services dhcp release *interface-name*** operational command had no impact. [PR1072319](#)

Software Installation and Upgrade

- In a mixed EX4200 and EX4500 Virtual Chassis or in an EX3300 Virtual Chassis, or on an EX6200 or EX8200 switch, during a nonstop software upgrade (NSSU), packets might be duplicated. [PR1062944](#)
- On an EX8200 Virtual Chassis, an NSSU to Junos OS Release 15.1R1 might fail after the image is pushed to the backup Routing Engine, and a vmcore might be created. [PR1075232](#)
- In Junos Space, the Junos OS Release 15.1R1 image for EX9200 switches is not mapped to the correct platform. As a workaround, in Junos Space, right-click the device image, and select **ex-92xx** in **Modify device image**. [PR1090863](#)
- On EX9200 switches, during an in-service software upgrade (ISSU) from Junos OS Release 15.1R1 to Release 15.1R2, BGP and Layer 3 multicast traffic might be dropped for approximately 30 seconds. [PR1116299](#)

Spanning-Tree Protocols

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

Virtual Chassis

- On an EX9200 Virtual Chassis, if you restart an FPC with Virtual Chassis ports (VCPs) and there are no other FPCs with VCPs, a Virtual Chassis split might occur and the backup FPC might show a machine check exception and create a Network Processing Card (NPC) core file. [PR1083965](#)
- On an EX9200 Virtual Chassis with JDHCP_Relay_LSYS configurations, the Virtual Chassis linecard members might go up and down after you reboot the switch. [PR1108402](#)

Related Documentation

- [New and Changed Features on page 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R3 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.

- [High Availability \(HA\) and Resiliency on page 46](#)
- [Infrastructure on page 47](#)
- [Port Security on page 47](#)
- [Software Installation and Upgrade on page 47](#)

High Availability (HA) and Resiliency

- Substantial traffic losses might occur during a nonstop software upgrade (NSSU) in a mixed EX4200 and EX4500 Virtual Chassis, in an EX3300 Virtual Chassis, on an EX6200 switch, on an EX8200 switch, or in an EX8200 Virtual Chassis. [PR1062960](#)
- On an EX4300 Virtual Chassis and on EX8200 switches, when you perform an NSSU, there might be up to five seconds of traffic loss for multicast traffic. [PR1125155](#)
- On EX4300 Virtual Chassis, NSSU is not supported from Junos OS Release 14.1X53-D35 to Release 15.1. [PR1148760](#)

Infrastructure

- On EX4300 switches, starting in Junos OS Release 15.1R3, a pfex_junos core file might be created when you add or delete a native VLAN configuration with **flexible-vlan-tagging**. [PR1089483](#)
- On EX2200 and EX3300 switches, ARP requests might be dropped when IP source guard is enabled and 802.1X (dot1x) authentication assigns a new dynamic VLAN to the client MAC. [PR1169150](#)

Port Security

- On EX3300 switches, ARP requests might be dropped when IP source guard is enabled and 802.1X (dot1x) authentication assigns a new dynamic VLAN to the client MAC. [PR1062960](#)

Software Installation and Upgrade

- Substantial traffic losses might occur during an NSSU upgrade on EX4200 and EX4500 Virtual Chassis, EX6200 and EX8200 switches, or EX8200 Virtual Chassis. [PR1062960](#)
- On EX9200 switches, unified ISSU does not work properly for upgrading to Junos OS Release 15.1R1. Junos Space triggers the upgrades and the upgrades fail. [PR1091610](#)
- On EX9200 switches, after an ISSU is performed, storm control takes effect only after you delete the storm control configuration and then re-create it. [PR1151346](#)

Related Documentation

- [New and Changed Features on page 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

Resolved Issues

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

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

- [Resolved Issues: Release 15.1R3 on page 48](#)
- [Resolved Issues: Release 15.1R2 on page 53](#)

Resolved Issues: Release 15.1R3



NOTE: Some resolved issues at Release 15.1R3 apply to both QFX Series and EX Series switches. Those shared issues are listed in the QFX Series “[Resolved Issues](#)” on page 243: Release 15.1R3 section.

- [Authentication and Access Control](#)
- [Dynamic Host Configuration Protocol](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Multicast](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Software Installation and Upgrade](#)
- [Spanning-Tree Protocols](#)
- [Virtual Chassis](#)

Authentication and Access Control

- On EX2200 switches, if you issue the CLI command **request system services dhcp release interface-name**, an IP address release message DHCP packet is sent from the client and processed at the server. At the same time, the client clears the IP address on the same interface, and the clearance of the IP address on the interface leads to the acquisition of a new IP address from the server. If you then issue the CLI command **show system services dhcp client interface-name**, the output of this operational command indicates that the command had no impact. [PR1072319](#)
- On an EX2200 or EX3300 switch on which Dynamic Host Configuration Protocol (DHCP) relay is enabled, when a client requests an IP address, the system might generate a harmless warning message such as: `/kernel: Unaligned memory access by pid 19514 [jdhcpd] at 46c906 PC[104de0]`. [PR1076494](#)
- On EX9200 switches, when 802.1X (dot1x) authentication is configured, the **show dot1x authentication-failed-users** command output might not show the Failure Count attribute correctly. [PR1080451](#)
- On EX Series switches, if 802.1X authentication (dot1x) is configured on all interfaces, an 802.1X-enabled interface might get stuck in the *Initialize* state after the interface goes down and comes back up, and 802.1X authentication fails. Also, if 802.1X authentication (dot1x) is configured on all interfaces and the **no-mac-table-binding** configuration statement is configured under the **[edit protocols dot1x authenticator]** hierarchy level, the dot1x process (dot1xd) might generate core files after it is deactivated and then reactivated, and 802.1X authentication might be temporarily impacted until the process restarts automatically. [PR1127566](#)

- On EX Series switches, the **use-option-82** statement under the **[edit ethernet-switching-options secure-access-port vlan *vlan-name* dhcpv6-option18]** hierarchy might not work as expected after you commit the configuration. [PR1146588](#)
- On EX4300 switches, if you change the server-fail VLAN, all authenticated supplicants are disconnected. They are then authenticated again, and during this disconnection and reconnection, there is a service impact for three through four seconds. [PR1151234](#)

Dynamic Host Configuration Protocol

- On EX9200 switches, DHCP snooping and related access security features ARP inspection, IP source guard, Neighbor Discovery inspection, and IPv6 source guard, are not supported at the **[edit logical-systems *logical-system-name* vlans *vlan-name* forwarding-options dhcp-security]** hierarchy level. [PR1087680](#)

High Availability (HA) and Resiliency

- On EX8200 switches, a nonstop software upgrade (NSSU) might fail during the master Routing Engine upgrade step, and an NSSU process might abort with this message: **mgd: unable to execute /var/etc/reboot.ex: Authentication error.** [PR1122628](#)

Infrastructure

- On EX2200 switches, system log messages might display IP addresses in reverse order. For example, an ICMP packet from 10.0.1.114 to 10.0.0.7 might be displayed in the log as: **PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 114.1.0.10 7.0.0.10 0 0 (1 packet).** The correct log message is: **PFE_FW_SYSLOG_IP: FW: ge-0/0/0.0 R icmp 10.0.1.114 10.0.0.7 0 0 (1 packet).** [PR898175](#)
- On EX2200 and EX3300 Virtual Chassis, the Internal state in ERPS is not updated properly in certain conditions. As a workaround, check the interface state and update the ERPS engine accordingly so that they are always in sync. [PR975104](#)
- On EX4300 switches, if a Gigabit Ethernet interface is directly connected to an MX104 management interface (fxp0), the physical link will be down. [PR1069198](#)
- On EX4300 switches, traffic sampling is not supported. If you configure traffic sampling, the sampling process (sampled) might generate a core file. [PR1091826](#)
- On an EX4300 Virtual Chassis or a mixed mode Virtual Chassis that has an EX4300 as a member, if you disable root login connections to the console port by issuing the **set system ports console insecure** command, users can still log in as root from the backup and linecard members of the Virtual Chassis. [PR1096018](#)
- On EX4600 switches, the EX4600-EM-8F expansion module interfaces might not come up if the module is removed and re-inserted or if the PIC is taken offline and then brought online. [PR1100470](#)
- On EX8200 switches with multicast protocols configured, when a multicast-related (non-aggregated Ethernet) interface goes down and comes back up, ARP installation for certain hosts might fail because stale entries have not been cleared, and traffic might be lost as well. [PR1105025](#)
- On EX4200 switches with multiple member interfaces on an aggregated Ethernet (AE) interface and with a large-scale CoS configuration enabled on the AE

interface, a Packet Forwarding Engine limit might be exceeded, the Packet Forwarding Engine might return an invalid ID, and the Packet Forwarding Engine manager (pfem) process might generate core files. [PR1109022](#)

- On EX4500 or EX4550 Virtual Chassis, if an NFS/UDP fragmented packet enters the Virtual Chassis through a LAG and traverses a Virtual Chassis port (VCP) link, CPU utilization might become high, and the software forwarding infrastructure (sfid) process might generate a core file. [PR1109312](#)
- On EX Series switches, an interface with an EX-SFP-1GE-LH transceiver might not come up and the transceiver might be detected as an SFP-EX transceiver. [PR1109377](#)
- On EX9200 switches, starting with Junos OS Release 14.1R1, 32k is the minimum value that you must configure for policer bandwidth limits. If you configure a policer bandwidth limit that is less than 32k, an error message is displayed. [PR1109780](#)
- On EX4500 switches, if MPLS and CoS behavior aggregate (BA) classifiers are configured on the same interface, the BA classifiers might not work. As a workaround, use multifield (MF) classifiers instead of BA classifiers. [PR1116462](#)
- On EX4200 and EX4550 switches, the xe- interfaces in a 10-gigabit SFP+ expansion module (EX4550-EM-8XSFP) or an SFP+ MACsec uplink module (EX-UM-2X4SFP-M) might stop forwarding traffic if the module is removed and reinserted or if the PIC goes offline and comes back online. [PR1113375](#)
- On EX Series switches, if you deactivate an output interface that is configured with **family mpls**, a nondefault CoS classifier configured on the interface might be deleted, placing traffic in the wrong queue. [PR1123191](#)
- On EX4300 switches, when there is a redundant trunk group (RTG) link failover, media access control (MAC) refresh packets might be sent out from a non-RTG interface that is in the same VLAN as the RTG interface, and a traffic drop might occur because of MAC flapping. [PR1133431](#)
- On EX9200 switches, the Layer 2 address learning daemon (l2ald) might crash continuously and create core files after you configure the fxp0 interface as **ethernet-switching** and commit the configuration. [PR1127324](#)
- On EX4300 switches, if the switch works as part of a target subnet, while receiving the targeted broadcast traffic, packets might be forwarded to the destination with the switch's MAC address as the destination MAC address, rather than the Layer 2 broadcast frame with destination MAC address FFFF.FFFF.FFFF. [PR1127852](#)
- On EX Series switches, an interface with a non-Juniper Networks 1000BASE-EX SFP Module-40km might not come up because register values are not set to correct values. This issue occurs only during initial deployment of the switch or when the switch is upgraded to Junos OS Release 12.3R8, 13.2X51-D30, 14.1X53-D10, or 15.1R2 onwards. [PR1142175](#)
- On EX9200 switches, an IRB unicast next hop in a scenario with a Layer 2 LAG as the underlying interface might result in traffic blackholing. [PR1114540](#)
- On EX9200 switches, a secondary VLAN might be mapped to the primary VLAN IRB interface to facilitate ARP synchronization across MC-LAG peers running a PVLAN configuration. [PR1145623](#)

Interfaces and Chassis

- If an EX4550-32F switch in a Virtual Chassis reboots and comes online, LACP interfaces on any of the member switches of the Virtual Chassis might go down and not come up. [PR1035280](#)
- On a two-member EX8200 Virtual Chassis, if the Link Aggregation Control Protocol (LACP) child interfaces span different Virtual Chassis members, the MUX state in the LAG member interfaces might remain in the *Attached* or *Detached* state after you disable and then reenables the AE interface. [PR1102866](#)

Layer 2 Features

- On EX Series switches, if you configure Ethernet ring protection (ERP) with interfaces configured with **vlan members all**, commit the changes, then add a new VLAN and commit the configuration again, the Ethernet switching process (eswd) might crash when a non-ERP interface goes down and then comes back up. [PR1129309](#)
- On EX Series switches except EX4300, EX4600, and EX9200, the Ethernet switching process (eswd) might crash if you delete a VLAN tag and then add the VLAN name by using a single commit, in the configuration under the **[edit ethernet-switching-options unknown-unicast-forwarding]** hierarchy. [PR1152343](#)

Multicast

- On EX Series switches, unregistered multicast packets are not filtered and are instead forwarded to all unexpected ports, even though IGMP snooping is enabled. [PR1115300](#)
- On an EX3300 switch, if you configure IGMP snooping with a VLAN that is not on the switch, the commit fails. [PR1149509](#)

Network Management and Monitoring

- On EX Series switches (except EX4300, EX4600, and EX9200), when system log is enabled and an RPM probe is set to greater than 8000 bytes, the message **?PING_RTT_THRESHOLD_EXCEEDED?** is not displayed, although it should be. [PR1072059](#)
- On EX Series switches, there are two issues regarding SNMP MIB walks: A private interface—for example, pime.32769—must have an ifIndex value of less than 500. If you do not add the private interface to a static list of rendezvous point (RP) addresses, the mib2d process assigns an ifIndex value from the public pool (with ifIndex values greater than 500) to the interface, which then will have an incorrect ifIndex allocation. A random **Request failed: OID not increasing** error might occur when you issue the **show snmp mib walk** command, because the kernel response for a 10-gigabit interface during an SNMP walk might take more than one second, and the mib2d process receives duplicate SNMP queries from the snmpd process. [PR1121625](#)
- On EX9200 switches, the value for the **udpOutDatagrams** object displayed in the output of the **show snmp mib walk decimal udpOutDatagrams** command is different from that displayed for the same object in the output of the **show system statistics udp member 0** command. The value for the **datagrams dropped due to no socket** field is incorrectly used as the **udpOutDatagrams** value in the output for **show snmp mib walk decimal**

udpOutDatagrams. As a workaround, use the **show system statistics udp member 0** command. [PR1104831](#)

Platform and Infrastructure

- On EX4300 switches with redundant trunk groups (RTGs) configured, after an RTG primary link comes online from the offline state, it becomes the active link and the other link becomes the backup link. After this, the Layer 2 address learning daemon (l2ald) sends a MAC refresh packet out of the new active RTG logical interface, which is not yet programmed in the Packet Forwarding Engine. This causes the primary link to incorrectly update the MAC entry and also causes traffic loss. [PR1095133](#)
- On EX4300 switches with Virtual Router Redundancy Protocol (VRRP) configured on an integrated routing and bridging (IRB) logical interface, when the IRB logical interface is disabled or deleted, the kernel does not send VRRP dest-mac-filter delete messages to the Packet Forwarding Engine, which might cause loss of traffic that comes from another device's same VRRP group master VIP to the backup (or backup to master). [PR1103265](#)
- On EX4300 switches, VSTP BPDUS are not flooded in the VLAN when VSTP is not configured on the switches. [PR1104488](#)
- On EX4300 switches, if a policer ICMP filter is applied on the loopback interface, incoming ICMP packets might be dropped on the ingress Packet Forwarding Engine and ARP requests might not be generated. [PR1121067](#)
- On EX4300 switches, configuring **set groups group_name interfaces interface-name unit 0 family ethernet-switching** and committing the configuration might cause the Layer 2 address learning process (l2ald) to generate a core file. [PR1121406](#)
- On EX4300 switches, port vector corruption on a physical port might be caused by the interface flapping multiple times, which leads to a Packet Forwarding Engine manager (pfem) crash and a Routing Engine reboot. [PR1121493](#)
- On EX4300 switches with a Q-in-Q configuration, when Layer 2 protocol tunneling (L2PT) for VLAN Spanning Tree Protocol (VSTP) is enabled, the C-VLAN (inner VLAN or customer VLAN) might not be encapsulated in the PDUs that exit the trunk port. [PR1121737](#)
- On an EX4300 Virtual Chassis, if a redundant trunk group (RTG) interface flaps, when control packets originating from the switch are going over that RTG interface, the core device become nonresponsive and you would have to reload the device to restore connectivity. [PR1130419](#)
- On EX4300 Virtual Chassis, traffic from or to a Routing Engine through an aggregated Ethernet (AE) member interface that is not in the master might be dropped, but traffic transmitted through the switch (that is, hardware switched) is not affected. [PR1130975](#)
- On an EX4300 switch, when an SNMP walk is performed to query the native VLAN, for most of the trunk interfaces, the query might return a value of 0 instead of the configured native VLAN ID. [PR1132752](#)
- On EX4300 switches configured with Ethernet ring protection switching (ERPS), the ping might not go through after the Wait to Restore (WTR) timer expires. [PR1132770](#)

- On EX4300 switches, a filter might not work as expected when you commit a filter-based forwarding (FBF) configuration for the first time after rebooting the switch. [PR1135771](#)
- On EX Series switches, the following DEBUG messages might be incorrectly displayed as output with logging level INFO: %USER-6: [EX-BCM PIC] ex_bcm_pic_eth_an_config %USER-6: [EX-BCM PIC] ex_bcm_pic_check_an_config_change [PR1143904](#)
- On EX4300 switches, if an IPv6 firewall filter term exceeds the maximum, the Packet Forwarding Engine manager (pfex) might crash continuously. [PR1145432](#)
- On EX4300 switches with redundant trunk groups (RTGs) configured, VSTP BPDUs coming into an RTG backup interface might be incorrectly forwarded out of interfaces other than the RTG primary interface. [PR1151113](#)

Software Installation and Upgrade

- On EX8200 switches, an NSSU from Junos OS Release 15.1R1 to Release 15.1R2 fails with the message: **mgd: unable to execute /var/etc/reboot.ex: Authentication error.** [PR1122628](#)

Spanning-Tree Protocols

- On EX Series switches with dual Routing Engines or on an EX Series Virtual Chassis, the switch or the Virtual Chassis might send multiple proposal BPDUs on an alternate port after a Routing Engine switchover or a nonstop software upgrade (NSSU), resulting in the peer device receiving multiple proposal BPDUs and triggering a dispute condition. The peer port states constantly alternate between *FORWARDING* and *BLOCKING*. [PR1126677](#)
- On EX Series switches with bridge protocol data unit (BPDU) protection configured on all edge ports, edge ports might not work correctly and might revert to the unblocking state when the **drop** option is configured under the **[edit ethernet-switching-options bpdudrop interface xstp-disabled]** hierarchy. [PR1128258](#)

Virtual Chassis

- On a two-member EX Series Virtual Chassis in which the same mastership priority is configured on both members, if there are more than 34 SFPs present in the current master and if a reboot is issued in the current master, then the backup becomes the master. When the original master rejoins the Virtual Chassis, it regains mastership. [PR1111669](#)

Resolved Issues: Release 15.1R2

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

- [Spanning-Tree Protocols](#)
- [VPLS](#)

Class of Service (CoS)

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

Dynamic Host Configuration Protocol

- On EX9200 switches, when DHCP relay is configured using the **forward-only** and **forward-only-replies** statements at the **[edit forwarding-options dhcp-relay]** hierarchy level, if the DHCP local server is also configured with the **forward-snooped-clients** statement at the **[edit system services dhcp-local-server]** hierarchy level, the configuration for **forward-snooped-clients** takes precedence over the configuration for **forward-only** and **forward-only-replies**. As a result, DHCP message exchange between VRFs might not work as expected. [PR1077016](#)
- On EX Series switches except EX9200, the configuration of options for the **circuit-id** CLI statement at the **[edit forwarding-options dhcp-relay group group-name relay-option-82]** hierarchy level does not work as expected. The format of the DHCP option 82 Circuit ID must be **switch-name:physical-interface-name:vlan-name**, but instead, the format is **switch-name:vlan-name**. [PR1081246](#)
- On EX Series switches except EX9200 switches, with DHCP relay configured on the IRB interface for BOOTP relay, if the client is connected to the physical interface that belongs to the same VLAN as the IRB interface, and sends BOOTP request packets to the server, BOOTP reply packets from the server might be dropped on the IRB interface. [PR1096560](#)

Interfaces and Chassis

- On EX9200 switches, if an interface range is configured that includes large-scale physical interfaces, and with the **family** option set to **ethernet-switching**, the configuration might take a long time to commit. [PR1072147](#)
- On EX9200 switches, if an interface for which the MAC move limit action is set to **shutdown** goes down and comes up, and then a Layer 2 learning (l2ald) process restarts, the logical interface remains down even if you issue the command **clear ethernet-switching recovery-timeout**. [PR1072358](#)
- On EX9200 switches, when a MAC move limit is configured on two VLAN members and the limit is configured with the action **vlan-member-shutdown** on two VLAN members, if the limit is reached on one VLAN member, both members are disabled, blocking all traffic. [PR1078676](#)
- On EX9200 platforms, if you configure an MC-LAG with two devices, and then delete and re-create an MC-AE interface, broadcast and multicast traffic that is flooded might loop for several milliseconds. [PR1082775](#)
- An EX9200-40F-M line card drops all traffic on an IRB logical interface, including both data plane and control plane traffic. If an IRB logical interface is configured on an EX9200-40F-M line card as part of a VLAN, any device connected through that interface

cannot use Layer 3 forwarding outside the subnet, because the EX9200-40F-M line card does not handle the ARP function correctly. Configuring static ARP on devices using the EX9200 as a gateway is not a workaround, because packets are still dropped if the Routing Engine of the EX9200 has the routes and ARP entry for the destination IP. [PR1086790](#)

Media Access Control Security (MACsec)

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

Network Management and Monitoring

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

Platform and Infrastructure

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

Spanning-Tree Protocols

- On EX Series Virtual Chassis, if STP is configured, and each member's mastership priority values are different, rebooting some or all of the Virtual Chassis members might cause a traffic failure, even after the reboot has completed. [PR1066897](#)

- On EX Series switches except EX9200, when MSTP is configured, the Ethernet switching process (eswd) might generate multiple types of core files in the large-scale VLANs that are associated with multiple spanning-tree Instances (MSTIs). [PR1083395](#)

VPLS

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

Related Documentation

- [New and Changed Features on page 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Documentation Updates on page 56](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

Documentation Updates

There are no errata or changes in Junos OS Release 15.1R3 for the EX Series switches documentation.

Related Documentation

- [New and Changed Features on page 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)
- [Product Compatibility on page 57](#)

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 57](#)

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths—you can upgrade directly from one EEOL release to the next EEOL release, even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 10.0, 10.4, and 11.4 are EEOL releases. You can upgrade from Junos OS Release 10.0 to Release 10.4 or even from Junos OS Release 10.0 to Release 11.4. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 10.3 (a non-EEOL release) to Junos OS Release 11.4 or directly downgrade from Junos OS Release 11.4 to Junos OS Release 10.3.

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

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

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

Related Documentation

- [New and Changed Features on page 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)
- [Product Compatibility on page 57](#)

Product Compatibility

- [Hardware Compatibility on page 57](#)

Hardware Compatibility

To obtain information about the components that are supported on the devices, and 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 36](#)
- [Changes in Behavior and Syntax on page 43](#)
- [Known Behavior on page 43](#)
- [Known Issues on page 46](#)
- [Resolved Issues on page 47](#)
- [Documentation Updates on page 56](#)
- [Migration, Upgrade, and Downgrade Instructions on page 56](#)

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 15.1R3 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: For M Series, MX Series, and T Series routers, unified ISSU upgrade from Junos OS release 15.1R3 to Junos OS Release 15.1F4 is not supported.

- [New and Changed Features on page 59](#)
- [Changes in Behavior and Syntax on page 100](#)
- [Known Behavior on page 125](#)
- [Known Issues on page 128](#)
- [Resolved Issues on page 142](#)
- [Documentation Updates on page 200](#)
- [Migration, Upgrade, and Downgrade Instructions on page 205](#)
- [Product Compatibility on page 215](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R3 for the M Series, MX Series, and T Series.

- [Hardware on page 60](#)
- [Bridging and Learning on page 61](#)
- [Class of Service \(CoS\) on page 61](#)
- [High Availability \(HA\) and Resiliency on page 62](#)
- [Interfaces and Chassis on page 64](#)
- [IPv6 on page 69](#)
- [Junos OS XML API and Scripting on page 69](#)
- [Layer 2 Features on page 69](#)
- [Management on page 71](#)
- [MPLS on page 71](#)
- [Multicast on page 73](#)
- [Network Management and Monitoring on page 75](#)
- [Routing Policy and Firewall Filters on page 76](#)
- [Routing Protocols on page 77](#)

- [Services Applications on page 80](#)
- [Software Defined Networking on page 84](#)
- [Software Installation and Upgrade on page 84](#)
- [Subscriber Management and Services \(MX Series\) on page 85](#)
- [User Interface and Configuration on page 98](#)
- [VPNs on page 98](#)

Hardware

- **New MPC variants that support higher scale and bandwidth (MX Series)**—Starting with Junos OS Release 15.1, the following variants of a new MPC with higher scale and bandwidth are supported on MX Series routers:
 - MPC2E-3D-NG—80 Gbps capacity without hierarchical quality of service (HQoS)
 - MPC2E-3D-NG-Q—80 Gbps capacity with HQoS
 - MPC3E-3D-NG—130 Gbps capacity without HQoS
 - MPC3E-3D-NG-Q—130 Gbps capacity with HQoS

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



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



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

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

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

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

Bridging and Learning

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

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

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

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



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

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

Class of Service (CoS)

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

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

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

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

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

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

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

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



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

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

High Availability (HA) and Resiliency

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

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

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

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

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

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

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

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

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

- RELAYD_COMMAND_OPTIONS
- RELAYD_COMMAND_OPTION_ERROR
- RELAYD_SYSCALL_ERROR
- **Configuration support for multiple MEPs for interfaces belonging to a single VPLS service, CCC, or bridge domain (MX Series)**—Starting with Junos OS Release 15.1, you can configure multiple maintenance endpoints (MEPs) for a single combination of maintenance association and maintenance domain IDs for interfaces belonging to a particular VPLS service, circuit cross-connect (CCC), or bridge domain.

To configure multiple MEPs, use the existing **mep mep-id** statement at the **[edit protocols oam ethernet connectivity-fault-management maintenance-domain domain-name maintenance association ma-name]** hierarchy level.
- **NSR and validation-extension for BGP flowspec**—Starting in Junos OS Release 15.1, changes are implemented to add NSR support for existing inet-flow and inetvpnflow families and to extend routes validation for BGP flowspec. Two new statements are introduced as part of this enhancement.

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

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

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

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

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

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

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

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

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

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

Interfaces and Chassis

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

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

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

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

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

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

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

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

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

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

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

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

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

connectivity-association-name statement and its substatements at the **[edit security macsec]** hierarchy level. MACsec is supported on MX Series routers with MACsec-capable interfaces. MACsec using static secure association key (SAK) security mode does not work properly on MX80 routers and FPC slots other than slot 0 of MX104 routers.

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

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

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

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

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

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

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

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

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

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

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

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

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

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



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

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

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

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

- **QSFP-40GE-LX4 (MX Series)**—In Junos OS Release 15.1R3 and later, the QSFP-40GE-LX4 transceiver provides 2km reach over single-mode fiber, 100m (with OM3 MMF cable), or 150m (with OM4 MMF cable) reach over multimode fiber. Signaling speed for each channel is 10.3125 GBd with aggregated data rate 41.25 Gb/s. The module enables 40GBASE links over a pair of either SMF or MMF terminated with duplex LC connectors. The LC connector supports connections with physical contact (PC) or ultra physical contact (UPC) connectors. Patch cords with APC connectors are not supported. The 6x40GE +24X10GE MPC5EQ (model number: MPC5EQ-40G10G) supports the QSFP-40GE-LX4 transceiver.

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

[See [40-Gigabit Ethernet 40GBASE-R Optical Interface Specifications](#).]

- **CFP2-100G-ER4-D (MX Series)**—In Junos OS Releases 13.3R9, 14.2R6, and 15.1R3 and later, the CFP2-100G-ER4-D transceiver provides dual-rate 40 km reach over G.652 single-mode fiber. Signaling speed for each channel is either 25.78125 GBd with aggregated data rate 103.125 Gb/s of 100GBASE-R, or 27.952493 GBd with aggregated data rate 111.81 Gb/s of OTU4 client interface. The CFP2-100G-ER4-D transceiver supports both IEEE 100GBASE-ER4 and ITU-T G.959.1 application code 4L1-9C1F. The duplex LC connector supports connections with Physical Contact (PC) or Ultra Physical Contact (UPC) connectors. Patch cords with APC connectors are not supported. The CFP2-100G-ER4-D supports the 100GBASE-ER4 standard. The following MPCs and MIC support the CFP2-100G-ER4-D transceiver:
 - 2x100GE + 4x10GE MPC5E (model number: MPC5E-100G10G)
 - 2x100GE + 4x10GE MPC5EQ (model number: MPC5EQ-100G10G)
 - 100-Gigabit Ethernet MIC with CFP2 (model number: MIC6-100G-CFP2)

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

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#).]

- **CFP2-100G-SR10-D3 (MX Series)**—In Junos OS Release 15.1R3 and later, the CFP2-100G-SR10-D3 transceiver provides dual rate 100 m (with OM3 MMF cable) and 150 m (with OM4 MFF cable) reach over multimode fiber. Signaling speed for each channel is either 10.3125 GBd with aggregated data rate 103.125 Gb/s of 100GBASE-R, or 11.181 GBd with aggregated data rate 111.81 Gb/s of OTU4 client interface. With 24-fiber ribbon cables that have MPO connectors, the module can support 100-gigabit links. With ribbon to duplex fiber breakout cables, the module can also support 10 x 10 Gigabit mode. The recommended Option A in IEEE STD 802.3-2012 is required. The CFP2-100G-SR10-D3 transceiver supports the 100GBASE-SR10 standard. The following MPCs and MIC support the CFP2-100G-SR10-D3 transceiver:
 - 2x100GE + 4x10GE MPC5E (model number: MPC5E-100G10G)
 - 2x100GE + 4x10GE MPC5EQ (model number: MPC5EQ-100G10G)
 - 100-Gigabit Ethernet MIC with CFP2 (model number: MIC6-100G-CFP2)

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

[See [100-Gigabit Ethernet 100GBASE-R Optical Interface Specifications](#).]

IPv6

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

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

Junos OS XML API and Scripting

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

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

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

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

Layer 2 Features

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

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

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

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

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

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

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

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

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

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

Management

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

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

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

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

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

MPLS

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

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

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

You can configure pro-active loss and delay measurement using the **performance-monitoring** configuration statement. This functionality provides real-time

visibility into network performance, thereby facilitating network performance planning, troubleshooting, and evaluation.

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

- **Configuring Layer 3 VPN egress protection with PLR as protector (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, this feature addresses a special scenario of egress node protection, where the point of local repair (PLR) and the protector are co-located as one router. In this case, there is no need to have a bypass LSP reroute traffic during local repair.

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

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

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

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

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

- **Support for DDoS on pseudowire subscriber logical interface (M Series and MX Series)**—Starting with Junos OS Release 15.1R3, distributed denial-of-service (DDoS) protection is supported on the services side of an MPLS pseudowire subscriber logical interface. DDoS protection identifies and suppresses malicious control packets while enabling legitimate control traffic to be processed. This protection enables the device to continue functioning, even when attacked from multiple sources. Junos OS DDoS protection provides a single point of protection management that enables network administrators to customize a profile appropriate for the control traffic on their networks.
- **Support for Policer and Filter on pseudowire subscriber logical interface (M Series and MX Series)**—Starting with Junos OS Release 15.1R3, Policer and Filter are supported on the services side of an MPLS pseudowire subscriber logical interface. Policer defines a set of traffic rate limits and sets consequences for traffic that does not conform to the configured limits. Firewall filters restrict traffic destined for the Routing Engine based on its source, protocol, and application. Also, firewall filters limit the traffic rate of packets destined for the Routing Engine to protect against flood or denial-of-service (DoS) attacks.
- **Support for accurate transmit logical interface statistics on pseudowire subscriber logical interface (M Series and MX Series)**—Starting with Junos OS Release 15.1R3, accurate transmit logical interface statistics are supported on the services side of an MPLS pseudowire subscriber logical interface. These statistics report actual transmit

statistics instead of the load statistics given by the router for the pseudowire subscriber service logical interfaces.

- **Support for Ethernet circuit cross-connect (CCC) encapsulation on pseudowire subscriber logical interface (M Series and MX Series)**—Starting with Junos OS Release 15.1R3, CCC encapsulation is supported on the transport side of an MPLS pseudowire subscriber logical interface. This feature helps in migrating or deploying seamless MPLS architectures in access networks. Customers deploying either business edge or broadband residential edge access networks use this feature to configure interfaces over the virtual Ethernet interface similar to what is already available on physical Ethernet interfaces.

You can define only one transport logical interface per pseudowire subscriber logical interface. Although the unit number can be any valid value, we recommend that unit 0 represent the transport logical interface. Two types of pseudowire signaling are allowed, Layer 2 circuit and Layer 2 VPN.

- **MPLS over dynamic GRE tunnel scaling of 32K (MX Series)**—Starting in Junos OS Release 15.1R3, MX Series routers support dynamic GRE tunnels scaling to 32K. Additionally, the previous IFL dependency is removed so `rpdp` now creates a new tunnel composite nexthop rather than creating an IFL. The tunnel composite nexthop has encapsulation data of the dynamic tunnel with a VPN label. To enable nexthop base dynamic tunnel mode, you set the **next-hop-based-tunnel** statement from the **[routing-options]** hierarchy level. By configuring this new statement, you can switch an IFL-based tunnel to a nexthop-based dynamic tunnel. You can view output of this new statement with the following **show** commands: **show dynamic-tunnels database**, **show route table inet.3 extensive**, **show route table inet.3**, **show route table bgp.l3vpn.0**, and **show route table bgp.l3vpn.0 extensive**.



NOTE: Dynamic tunnels are not supported on logical systems.

Multicast

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

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

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

The default IGMP snooping implementation for a VPLS instance adds each pseudowire interface to its **oif** list. This includes traffic sent from the ingress PE router to the egress

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

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

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

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

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

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

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

Network Management and Monitoring

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

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

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

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

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

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

Junos OS also supports the **jnxDomLaneNotifications** traps.

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

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

A new enterprise-specific MIB, SOAM PM MIB, that defines the management objects for Ethernet services operations, administration, and maintenance for performance

monitoring, has been added and SNMP support is available for the MIB objects defined in Technical Specification MEF 36.

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

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

- **SNMP support for monitoring fabric power utilization (MX Series)**—Starting in Release 15.1, Junos OS supports monitoring of fabric power utilization. An SNMP trap is generated whenever the fabric power consumption exceeds the configured threshold value. The SNMP trap `jnxFabricHighPower` has been added to the `jnxFabricChassisTraps` group to indicate excessive power consumption. The SNMP trap `jnxFabricHighPowerCleared` added to the `jnxFabricChassisOKTraps` group sends notification when the condition of consuming excessive power is cleared.
- **Support for the interface-set SNMP index (MX Series)**—Starting with Release 15.1R2, Junos OS supports the interface-set SNMP index that provides information about interface-set queue statistics. The following interface-set SNMP index MIBs are introduced in the Juniper Networks enterprise-specific Class-of-Service MIB:
 - `jnxCosIfTable` in `jnxCos` MIB
 - `jnxCosIfsetQstatTable` in `jnxCos` MIB

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

Routing Policy and Firewall Filters

- **Support for consistent load balancing for ECMP groups (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, on MX Series routers with modular port concentrators (MPCs) only, you can prevent the reordering of flows to active paths in an ECMP group when one or more paths fail. Only flows that are inactive are redirected. This feature applies only to Layer 3 adjacencies learned through external BGP connections. It overrides the default behavior of disrupting all existing, including active, TCP connections when an active path fails. Include the **consistent-hash** statement at the **[edit policy-options policy-statement policy-statement-name then load-balance]** hierarchy level. You must also configure a global per-packet load-balancing policy.
[See [Actions in Routing Policy Terms](#).]
- **New fast-lookup-filter statement (MX240, MX480, MX960, MX2010, and MX2020 routers with MPC5E, MPC5EQ, and MPC6E MPCs and compatible MICs)**—Starting in Junos OS Release 15.1, the **fast-lookup-filter** option is available at the **[edit firewall family (inet | inet6) filter filter-name]** hierarchy level. This allows for hardware assist

from compatible MPCs in the firewall filter lookup. There are 4096 hardware filters available for this purpose, each of which can support up to 255 terms. Within the firewall, filters and their terms, ranges, prefix lists, and the except keyword are all supported. Only the inet and inet6 protocol families are supported.

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

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

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

- **Support for interfaces that use the same filter list to use a common template (MX5, MX10, MX40, and MX80 routers, and routers that use MX Series MPC line cards)**—Starting in Junos OS Release 15.1R3, on MX5, MX10, MX40, MX80, and MX Series routers with modular port concentrators (MPCs) only, you can configure all interfaces that use the same filter list to use a common template. This feature can be used to save microkernel memory and DME memory. Include the **filter-list-template** statement at the `[edit firewall family (inet | inet6) filter filter-name]` hierarchy level.

Routing Protocols

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

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

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

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

- **Multi-instance support for RSVP-TE (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 15.1, multi-instance support is extended to the existing MPLS

RSVP-Traffic Engineering (TE) functionality. This support is available only for a virtual router instance type. A router can create and participate in multiple independent TE topology partitions, which allows each partitioned TE domain to scale independently.

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

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

- **Support for long-lived BGP graceful restart (M Series, MX Series, and T Series)**—Starting in Release 15.1, Junos OS supports the mechanism to preserve BGP routing details from a failed BGP peer for a longer period than the duration for which such routing information is maintained using the BGP graceful restart functionality. To enable the BGP long-lived graceful restart capability, include the **long-lived receiver enable** statement at the **[edit protocols bgp graceful-restart]**, **[edit protocols bgp group group-name graceful-restart]**, and **[edit protocols bgp group group-name neighbor neighbor-address graceful-restart]** hierarchy levels.
 - **Selection of backup LFA for OSPF routing protocol (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1, the default loop-free alternative (LFA) selection algorithm or criteria can be overridden with an LFA policy. These policies are configured per destination per primary next-hop interface or per destination. These backup policies enforce LFA selection based on admin-group, srlg, node, bandwidth, protection-type, and metric attributes of the backup path. During backup shortest-path-first (SPF) computation, each attribute (both node and link) of the backup path, stored per backup next hop, is accumulated by IGP. For the routes created internally by IGP, the attribute set of every backup path is evaluated against the policy configured per destination per prefix primary next-hop interface. The first or the best backup path is selected and installed as the backup next hop in the routing table.
- [See [Example-configuring-backup-selection-policy-for-ospf-protocol](#).]
- **Remote LFA support for LDP in OSPF (MX Series)**—Beginning with Junos OS Release 15.1, you can configure a remote loop-free alternate (LFA) to extend the backup provided by the LFA in an OSPF network. This feature is useful especially for Layer 1 metro-rings where the remote LFA is not directly connected to the PLR. The existing LDP implemented for the MPLS tunnel setup can be reused for the protection of OSPF networks and subsequent LDP destinations thereby eliminating the need for RSVP-TE backup tunnels for backup coverage.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

or

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



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

Services Applications

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

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

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

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

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

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

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

- Configure a limit for inbound sessions for an EIF mapping. To configure this limit, include the **EIF-flow-limit** statement at the **[edit services nat rule rule-name term term-name then translated secure-nat-mapping]** hierarchy level.
- Configure a limit for the number of dropped flows (ingress, egress, or both) for a specified service set. To configure this limit, include the **max-drop-flows** statement at the **[edit services service-set service-set-name]** hierarchy level.

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

- **Support for per-service throughput for NAT and inline flow monitoring services (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1, you can configure the capability to transmit the throughput details per service for Junos Address Aware (carrier-grade NAT) and Junos Traffic Vision (previously known as J-Flow) in the last time interval to an external log collector. The default time interval at which the throughput data is sent is 300 seconds, which you can configure to suit your network needs. Multiple instances of the same service running on different PICs within a router are supported. This functionality is supported on MX Series routers with MS-MPCs and MS-MICs, and also in the MX Series Virtual Chassis configuration.
- **Support for generation of SNMP traps and alarms for inline video monitoring (MX Series)**—Starting in Junos OS Release 15.1, SNMP support is introduced for the media delivery index (MDI) metrics of inline video monitoring. Inline video monitoring is available on MX Series routers using only MPCE1, MPCE2, and MPC-16XGE. Until Junos OS Release 14.2, inline MDI generated only syslogs when the computed MDI metric value was not within the configured range. SNMP support is now added to enable SNMP traps to be triggered when the computed delay factor, media rate variation (MRV), or media loss rate (MLR) values are not within the configured range. You can retrieve the MDI statistics, flow levels, error details, and MDI record-level information using SNMP Get and Get Next requests. The SNMP traps and alarms that are generated when the MDI metrics exceed the configured ranges can be cleared as necessary. Also, you can control the flooding of SNMP traps on the system.
- **Support for Layer 2 services over GRE tunnel interfaces (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, you can configure Layer 2 Ethernet services over GRE interfaces (**gr-fpc/pic/port** to use GRE encapsulation). To enable Layer 2 Ethernet packets to be terminated on GRE tunnels, you must configure the bridge domain protocol family on the gr- interfaces and associate the gr- interfaces with the bridge domain. You must configure the GRE interfaces as core-facing interfaces, and they must be access or trunk interfaces. To configure the bridge domain family on gr- interfaces, include the **family bridge** statement at the **[edit interfaces gr-fpc/pic/port unit logical-unit-number]** hierarchy level.
- **Support for stateless source IPv6 prefix translation (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, you can configure stateless translation of source address prefixes in IPv6 networks. This capability is supported on MX Series routers with MPCs where inline NAT is supported. To configure stateless network prefix translation for IPv6 packets (NPTv6), include the **translation-type nptv6** statement at the **[edit services nat rule rule-name term term-name then translated]** hierarchy level. The NPTv6 translator translates the source address prefix in such a way that the transport layer checksum of the packet does not need to be recomputed.

- **Support for logging flow monitoring records with version 9 and IPFIX templates for NAT events (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1, you can configure MX Series routers with MS-MPCs and MS-MICs to log NAT events by using Junos Traffic Vision (previously known as J-Flow) version 9 or IPFIX (version 10) template format. NAT event logger generates messages in flow monitoring format for various NAT events, such as the creation of a NAT entry, deletion of a NAT entry, and for invalid NAT processing. These events also support NAT64 translations (translation of IPv6 addresses to IPv4 addresses), binding information base (BIB) events, and more detailed error generation. The generated records or logs for NAT events in flow template format are sent to the specified host or external device that functions as the NetFlow collector.
- **Support for unified ISSU on inline LSQ interfaces (MX Series)**—Starting in Junos OS Release 15.1, unified in-service software upgrade (ISSU) is supported on inline link services intelligent queuing (IQ) (lsq-) interfaces on MX Series routers. Unified ISSU enables an upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. The inline LSQ logical interface (*lsq-slot/pic/0*) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC.
- **Inline TWAMP requester support (MX Series)**—Starting in Junos OS Release 15.1, MX Series routers support the inline Two-Way Active Measurement Protocol (TWAMP) control-client and session-sender for transmission of TWAMP probes using IPv4 between the sender (control-client) and the receiver (session-sender or server). The control-client and session-sender reside on the same router. The TWAMP control-client can also work with a third-party server implementation.
- **Ethernet over generic routing encapsulation (GRE) and GRE key support for label blocks (MX Series)**—Starting in Junos OS Release 15.1, MX Series routers support the following in compliance with RFC 2890:
 - Adding a bridge family on general tunneling protocol
 - Switching functionality supporting connections to the traditional Layer 2 network and VPLS network
 - Routing functionality supporting integrated routing and bridging (IRB)
 - Configuring the GRE key and performing the **hash load balance** operation both at the **gre tunnel initiated** and **transit routers** hierarchies
 - Providing statistics for the GRE-L2 tunnel
- **Support for IRB in a P-VLAN bridge domain (MX Series)**—Starting in Junos OS Release 15.1, MX Series routers support IRB in a private VLAN (P-VLAN) bridge domain. All IP features such as IP multicast, IPv4, IPv6, and VRRP that work for IRB in a normal bridge domain also work for IRB in a P-VLAN bridge domain.
- **Enhancements to the RFC 2544-based benchmarking tests (MX104)**—Starting in Junos OS Release 15.1, MX104 routers support RFC 2544-based benchmarking tests for Ethernet transparent LAN (E-LAN) services configured using LDP-based VPLS and BGP-based VPLS. The RFC 2544 tests are performed to measure and demonstrate the service-level agreement (SLA) parameters before the E-LAN service is activated. The tests measure throughput, latency, frame-loss rate, and back-to-back frames.

RFC 2544 performance measurement testing for Layer 2 E-LAN services on MX104 routers supports UNI-to-UNI unicast traffic only. You can enable reflection at the VPLS user-to-network interface (UNI). The following features are also supported:

- RFC2544 signature check—Verifies the signature pattern in the RFC2544 packets, by default.
- MAC swap for pseudowire egress reflection—Swaps the MAC addresses for pseudowire reflection.
- Ether type filter for both pseudowire and Layer 2 reflection—Specifies the ether type used for reflection.
- **Support for PCP version 2 (MX Series)**—Starting in Release 15.1, Junos OS supports Port Control Protocol (PCP) version 2, defined by IETF RFC 6887. PCP version 2 uses the client once for authentication. Junos OS is able to decode and process version 2 and version 1 messages. There are no CLI changes for PCP version 2 support.

[See [Port Control Protocol Overview](#).]

- **Support for inline MLPPP interface bundles on Channelized E1/T1 Circuit Emulation MICs (MX80, MX104, MX240, MX480, and MX960)**—Starting in Junos OS Release 15.1, you can configure inline MLPPP interfaces on MX80, MX104, MX240, MX480, and MX960 routers with Channelized E1/T1 Circuit Emulation MICs. Inline Multilink PPP (MLPPP), Multilink Frame Relay (FRF.16), and Multilink Frame Relay End-to-End (FRF.15) for time-division multiplexing (TDM) WAN interfaces provide bundling services through the Packet Forwarding Engine without requiring a PIC or Dense Port Concentrator (DPC). The inline LSQ logical interface (lsq-slot/pic/0) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC. A maximum of up to eight inline MLPPP interface bundles are supported on Channelized E1/T1 Circuit Emulation MICs, similar to the support for inline MLPPP bundles on other MICs with which they are compatible.
- **Data plane inline support for 6rd and 6to4 tunnels connecting IPv6 clients to IPv4 networks (MX Series with MPC5E and MPC6E)**—Starting with Release 15.1R3, Junos OS supports inline 6rd and 6to4 on MPC5E and MPC6E line cards. In releases earlier than Junos OS Release 15.1R3, inline 6rd and 6to4 was supported on MPC3E line cards only.

[See [Configuring Inline 6rd](#).]

- **Support for inline LSQ logical interface**—Starting in Junos OS Release 15.1R3, MPC2E-3D-NG and MPC3E-3D-NG support inline LSQ logical interface when flexible queuing is enabled. The inline LSQ logical interface (referred to as lsq-) is a virtual service logical interface that resides on the Packet Forwarding Engine to provide Layer 2 bundling services that do not need a service PIC.
- **Support for inline MPLS Junos Traffic Vision with IPFIX and v9 (MX Series)**—Starting in Junos OS Release 15.1, support of the MX Series routers for the inline Junos Traffic Vision feature is extended to the MPLS family consisting of the IP Flow Information Export (IPFIX) protocol and flow monitoring version 9 (v9). Currently, the inline Junos

Traffic Vision feature is supported only on the MS-MIC and MS-MPC consisting of the IPv4, IPv6, and virtual private LAN service (VPLS) protocols.

- **Support for inline 6rd and 6to4 (MX Series routers with MPC5Es and MPC6Es)**—Starting in Junos OS Release 15.1R3, you can configure inline 6rd or 6to4 on an MPC5E and MPC6E. 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.

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

Software Defined Networking

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

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

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

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

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

Software Installation and Upgrade

- **Validate system software add against running configuration on remote host or routing engine**—Beginning with Junos OS Release 15.1R2, you can use the **validate-on-host *hostname*** and **validate-on-routing-engine *routing-engine*** options with the **request system software add *package-name*** command to verify a candidate software bundle against the running configuration on the specified remote host or Routing Engine.
- **Validate system software against running configuration on remote host**—Beginning with Junos OS Release 15.1R2, you can use the **on (host *host* <username *username*> | routing-engine *routing-engine*)** option with the **request system software validate *package-name*** command to verify candidate system software against the running configuration on the specified remote host or Routing Engine.

- **Support for FreeBSD 10 kernel for Junos OS (MX240, MX480, MX960, MX2010, and MX2020)**—Starting with Junos OS Release 15.1, on the MX240, MX480, MX960, MX2010, and MX2020 only, FreeBSD 10 is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. Because installation restructures the file system, logs and configurations are lost unless precautions are taken. There are now Junos OS and OAM volumes, which provide the ability to boot from the OAM volume upon failures. Some system commands display different output and a few others are deprecated.

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

Subscriber Management and Services (MX Series)



NOTE: Although present in the code, the subscriber management features are supported in Junos OS Release 15.1R3 only for an early field qualification. Full support is available in a later release. Documentation for subscriber management features is included in the Junos OS Release 15.1 documentation set.

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

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

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

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

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

**NOTE:**

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

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

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

- **MAC address option for the Calling-Station-ID attribute (MX Series)**—Starting in Junos OS Release 15.1, you can specify that the subscriber MAC address is included in the Calling-Station-ID RADIUS attribute (31) that is passed to the RADIUS server. To do so, include the `mac-address` option when you configure the `calling-station-id-format` statement at the `[edit access profile profile-name radius options]` hierarchy level.

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

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

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

- **Support for overriding L2TP result codes (MX Series)**—Starting in Junos OS Release 15.1, you can configure the LNS to override result codes 4 and 5 with result code 2 in Call-Disconnect-Notify (CDN) messages. These result codes indicate that the number of L2TP sessions have reached the configured maximum value and the LNS can support no more sessions. When the LAC receives the code, it fails over to another LNS to establish subsequent sessions. Some third-party LACs respond only to result code 2. Include the `override-result-code session-out-of-resource` statement at the `[edit access-profile access-profile-name client client-name l2tp]` hierarchy level. Issue the `show services l2tp detail | extensive` command to display whether the override is enabled.

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

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

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

- **DHCPv6 relay agent Remote-ID (option 37) based on DHCPv4 relay agent information option 82 (MX Series)**—Starting in Junos OS Release 15.1, DHCPv6 relay agent supports a Remote-ID option (option 37) that is based on the DHCPv4 relay agent information option (option 82). When you enable this feature in dual-stack environments, the DHCPv6 relay agent checks the DHCPv4 binding for the option 82 Remote-ID suboption (suboption 2) and uses that information as option 37 in the outgoing RELAY-FORW message. In addition, you can specify the action DHCPv6 relay

agent takes if the DHCPv4 binding does not include an option 82 suboption 2 value; either forward the Solicit message without option 37 or drop the message.

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

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

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

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

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

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

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

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

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

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

- **Support for enhanced subscriber management subscriber logical interfaces or interface sets over MPLS pseudowires for a CoS scheduler hierarchy (MX Series)**—Starting in Junos OS Release 15.1, you can enable a CoS scheduling hierarchy

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

- **Support for enhanced subscriber management subscriber logical interfaces or interface sets over underlying logical interfaces for a CoS scheduler hierarchy (MX Series)**—Starting in Junos OS Release 15.1, you can enable a CoS scheduling hierarchy for enhanced subscriber management subscriber logical interfaces or interface sets over underlying logical interfaces. In Junos OS Release 14.2 and earlier, an interface set can be either at Layer 2 or Layer 3 levels of the CoS three-level hierarchical scheduler. You can now enable an enhanced subscriber management subscriber logical interface, such as an underlying logical interface, to function as a Layer 2 node in a CoS hierarchical scheduler. To configure three-level hierarchical scheduling, include the **implicit-hierarchy** statement at the **[edit interfaces "\$junos-interface-ifd-name" hierarchical-scheduler]** or **[edit interfaces lt-device hierarchical-scheduler]** hierarchy level.
- **Enhanced subscriber management support for ACI-based PPPoE subscriber session lockout (MX Series)**—Starting in Junos OS Release 15.1, enhanced subscriber management supports identification and filtering of PPPoE subscriber sessions by either the agent circuit identifier (ACI) value or the unique media access control (MAC) source address. You can use this feature when you configure PPPoE subscriber session lockout on static or dynamic VLAN and static or dynamic VLAN demux underlying interfaces.

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

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

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

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

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

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

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

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

- **Additional methods to derive values for L2TP connect speeds (MX Series)**—Starting in Junos OS Release 15.1, several new ways are supported for determining the transmit and receive connect speeds that the LAC sends to the LNS:
 - The Juniper Networks VSAs, Tx-Connect-Speed (26-162) and Rx-Connect-Speed (26-163), can provide the values.
 - The Juniper Networks VSA, Tunnel-Tx-Speed-Method (26-94), can specify a method (source) for the LAC to derive the values.
 - You can configure the LAC to use the actual downstream traffic rate enforced by CoS for the transmit speed. The **actual** method requires the effective shaping rate to be enabled and does not provide a receive speed, which is determined by the fallback scheme.

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

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

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

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

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

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

- **Support for Extensible Subscriber Services Manager (MX Series)**—Starting in Release 15.1, Junos OS supports Extensible Subscriber Services Manager (ESSM), a background process that enables dynamic provisioning of business services.
- **Loopback address as source address on DHCP relay agent**—Starting in Junos OS Release 15.1, you can configure the DHCPv4 and DHCPv6 relay agent to use the relay agent loopback address as the source address in DHCP packets. In network

configurations where a firewall on the broadband network gateway (BNG) is between the DHCP relay agent and the DHCP server, only the BNG loopback address passes through the BNG firewall. In that case, DHCP unicast packets do not pass through and are discarded. You can use two new configuration statements to override the DHCP source address with the BNG loopback address so DHCP packets do not pass through the firewall.

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

Now, when the SDB_FRAMED_PROTOCOL attribute is equal to AUTHD_FRAMED_PROTOCOL_PPP, the value of SDB_USER_IP_MASK is set to 255.255.255.255 by default. This value is overridden by the Framed-IP-Netmask value, if present.

When the SDB_FRAMED_PROTOCOL attribute is equal to AUTHD_FRAMED_PROTOCOL_PPP, the **show subscribers** command now displays the actual value of Framed-IP-Netmask in the IP Netmask field. Otherwise, the field displays the default value of 255.255.255.255.

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

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

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

When you first enable DHCP snooping, all DHCP traffic is snooped by default and only DHCP packets associated with subscribers (or their creation) will be handled, all other DHCP packets will be discarded. You can optionally modify this dropping behavior based on the type of interface—configured interfaces, non-configured interfaces, or all interfaces. All snooped DHCP traffic is still forwarded to the routing plane in the

routing instance, and in some cases, this results in excessive DHCP traffic being sent to the routing plane for snooping. The **no-snoop** statement disables snooping filters for DHCP traffic that can be forwarded directly from the hardware control plane, such as Layer 3 unicast traffic with a valid route, preventing that DHCP traffic from being forwarded to the slower routing plane of the routing instance.

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

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

- **Enhanced subscriber management support for source class usage in firewall filters (MX Series with MPCs)**—Starting in Junos OS Release 15.1R3, you can configure the **source-class** and **source-class-except** match conditions in a firewall filter that you create as part of a dynamic profile for use with enhanced subscriber management. Defining a firewall filter with matching based on source classes allows you to monitor the traffic of specific subscribers from specific network zones.

To configure a firewall filter term that matches an IPv4 or IPv6 source address field to one or more source classes, use the **source-class class-name** match condition at the **[edit dynamic-profiles profile-name firewall family family-name filter filter-name term term-name from]** hierarchy level. To configure a firewall filter term that does not match the IP source address field to the specified source classes, use the **source-class-except class-name** match condition at the same hierarchy level.

This feature enables you to dynamically configure firewall filters with the **source-class** and **source-class-except** match conditions as part of the same dynamic profile that activates services for a subscriber using enhanced subscriber management. In previous releases, you had to statically define the firewall filter outside of the dynamic profile used for service activation, which was a more time-consuming task and much less efficient.

- **Enhanced subscriber management support for configuring routing protocols in dynamic profiles (MX Series with MPCs)**—Starting in Junos OS Release 15.1R3, you can configure routing protocols (also known as routing services) on enhanced subscriber management interfaces as part of a dynamic profile. To do so, you must use the routing-services statement at the **[edit dynamic-profiles profile-name interfaces “\$junos-interface-ifd-name” unit “\$junos-underlying-interface-unit”]** hierarchy level.

When you enable enhanced subscriber management, the routing-services statement is required to configure all routing protocols except IGMP and MLD on dynamically created subscriber interfaces. The IGMP and MLD routing protocols are natively supported on enhanced subscriber management interfaces, and therefore do not require you to specify the routing-services statement.

When a dynamic profile containing the routing-services statement is instantiated, the router creates an enhanced subscriber management logical interface, also referred to as a pseudo logical interface, in the form **demux0.nnnn** (for example, **demux0.3221225472**). Any associated subscriber routes or routes learned from a routing protocol running on the enhanced subscriber management interface use this pseudo interface as the next-hop interface.

- **New commands for verifying and managing enhanced subscriber management (MX Series with MPCs)**—Starting in Junos OS Release 15.1R3, you can use the following operational commands to verify and manage enhanced subscriber management interfaces:
 - To display statistics information about enhanced subscriber management interfaces, use the **show system subscriber-management statistics** command. In addition to displaying basic packet statistics, you can use the available command options to view statistics specific to DHCP (dhcp), dynamic VLAN (dvlan), PPP (ppp), and PPPoE (pppoe) subscriber configurations.
 - To reset all statistics counters to zero, use the **clear system subscriber-management statistics** command.
 - To display information about how routes are mapped to specific enhanced subscriber management interfaces, use the **show system subscriber-management route** command. You can customize and filter the output by including one or more options in a single command.
- **Access Node Control Protocol agent support and limitations**—Starting in Junos OS Release 15.1R3, the Access Node Control Protocol (ANCP) agent requires enhanced subscriber management to be enabled, but support for the agent is limited to applying ANCP data to CoS traffic shaping for dynamic PPPoE and DHCP IP demux subscribers.

The ANCP agent does not support the following:

- Static or dynamic VLAN or VLAN demux interfaces.
- Static or dynamic interface-sets, including but not limited to agent circuit identifier (ACI) VLANs and VLAN-tagged interface-sets.
- RADIUS authentication or accounting.

- **Universal CAC for IPTV and VOD on MX Series Routers**—Starting in Junos OS Release 15.1R3, universal call admission control (CAC) is supported for multicast IPTV and unicast video on demand (VOD) traffic on MX Series routers. Universal CAC provides enhanced bandwidth management and prevents interface oversubscription to ensure high quality output by using dedicated and shared video bandwidth pools to limit the amount of traffic on subscriber interfaces.

To configure universal CAC, include the **access-cac** statement at the **[edit dynamic profiles profile name]** hierarchy level. You can then configure dedicated video bandwidth pools for IPTV by including the **multicast-video-bandwidth** statement, shared video bandwidth pools for IPTV and VOD by including the **video-bandwidth** statement, and multicast video policies by including the **multicast-video-policy** statement at the **[edit dynamic profiles profile name access-cac]** hierarchy level.

- **SNMP support for enhanced subscriber management dynamic interfaces**—Starting in Junos OS Release 15.1R3, SNMP support is available for enhanced subscriber management dynamic interfaces such as VLAN, PPP, and so on. An extension has been added to the Juniper Networks enterprise-specific Interface MIB to map enhanced subscriber management interfaces to logical route-mapping interfaces and to collect information about enhanced subscriber management interfaces. By default, data about enhanced subscriber management interfaces is not collected in the interfaces tables such as ifTable, ifXTable, and ifStackTable.

To enable querying of enhanced subscriber management interfaces through the Interface MIB, the Interface MIB must be configured at the interface level by enabling the **interface-mib** statement at the **[edit dynamic-profiles profile name interfaces interface-name]** hierarchy level. A link trap is sent for an enhanced subscriber management interface only if the interface name is present in ifTable and traps are enabled.

- **Enhanced subscriber management supported on the Multiservices Modular Interfaces Card (MS-MIC) and the Multiservices Modular PIC Concentrator (MS-MPC) (MX Series)**—Starting in Junos OS Release 15.1R3, the Carrier-Grade Network Address Translation (CGNAT) and inline flow monitoring services available with enhanced subscriber management support MS-MPCs and MS-MICs.
- **Captive portal content delivery (HTTP redirect) supported on the Multiservices Modular Interfaces Card (MS-MIC) and the Multiservices Modular PIC Concentrator (MS-MPC), and the Routing Engine (MX Series)**—Starting in Junos OS Release 15.1R3, 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.

- **Effective shaping rate and CoS adjustment control profiles on enhanced subscriber management interfaces (MX Series)**—Starting in Junos OS Release 15.1R3, CoS adjustment control profiles that determine the applications and algorithms that can modify a subscriber's shaping characteristics after a subscriber is instantiated are supported for enhanced subscriber management interfaces. Also, the effective shaping rate capability, which enables the actual downstream traffic rate to be computed and displayed, is also supported for enhanced subscriber management interfaces for accounting purposes.

When you configure CoS adjustment profiles and effective shaping rate on your router, the enhanced subscriber management interfaces that are defined as part of a dynamic profile at the **[edit dynamic-profiles profile-name interfaces “\$junos-interface-ifd-name” unit “\$junos-underlying-interface-unit”]** hierarchy level are considered for these functionalities. Only Ethernet interfaces are supported for these functionalities. Only dynamic subscribers are supported and static subscribers on enhanced subscriber management interfaces are not supported. Only the downstream shaping rate is validated and the upstream shaping rate is set to the advisory rate. Byte adjustments are not included in the effective shaping-rate. When cell-mode is specified, the Juniper Networks router adjusts rates (such as the shaping-rate) to “rate * 48/53” to account for 5-byte ATM AAL5 headers and does not account for cell padding.

- **Enhanced subscriber management support for ACI-based PPPoE subscriber session lockout (MX Series)**—Starting in Junos OS Release 15.1R3, 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.

- **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 15.1R3, 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 15.1R3, you can enable a CoS scheduling hierarchy for enhanced subscriber management subscriber logical interfaces or interface sets over underlying logical interfaces. In previous releases of Junos OS, an interface set could 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.
- **Changes in enhanced subscriber management support for allocating shared memory space (MX Series with MPCs)**—Starting in Junos OS Release 15.1R3, the first time you enable enhanced subscriber management, you must configure **max-db-size** for 400 MB or less (300MB is recommended). The **max-db-size** command can be found at the **[edit system configuration-database]** hierarchy level, and is used to allocated the amount of shared memory available to the configuration database.
- **Enhanced subscriber management on MX Series routers with MPCs**—Starting in Junos OS Release 15.1R3, you can configure and enable Junos OS enhanced subscriber management. Enhanced subscriber management is a next-generation broadband edge software architecture for wireline subscriber management. With enhanced subscriber management, you can take advantage of optimized scaling and performance for configuration and management of dynamic interfaces and services.

Configuring enhanced subscriber management consists of the following high-level tasks:

1. Download and install Junos OS Release 15.1R3, and reboot the router.



NOTE: Because unified in-service software upgrade (unified ISSU) is not supported when you upgrade to Junos OS Release 15.1R3, all subscriber sessions and subscriber state are lost after the upgrade.

2. Configure enhanced IP network services on the router.
3. Enable enhanced subscriber management.
4. Configure the maximum amount of shared memory (400 MB or less) used to store the configuration database for enhanced subscriber management.
5. (Optional) Enable graceful Routing Engine switchover (GRES) and nonstop active routing (NSR).
6. Commit the configuration and reboot the router.

After you configure and enable enhanced subscriber management, you can use dynamic profiles as usual for creating and managing dynamic subscriber interfaces and services.

- **Support for a static unnumbered interface with `$junos-routing-instance` (MX Series)**—Starting in Junos OS Release 15.1R3, you can configure a static logical interface as the unnumbered interface in a dynamic profile that includes dynamic routing instance assignment by means of the `$junos-routing-instance` predefined variable.



NOTE: This configuration fails commit if you also configure a preferred source address, either statically with the `preferred-source-address` statement or dynamically with the `$junos-preferred-source-address` predefined variable.



NOTE: The static interface must belong to the routing instance; otherwise the profile instantiation fails.

In earlier releases, when the dynamic profile includes the `$junos-routing-instance` predefined variable, you must do both of the following, else the commit fails:

- Use the `$junos-loopback-interface-address` predefined variable to dynamically assign an address to the unnumbered interface. You cannot configure a static interface address.
- Use the `$junos-preferred-source-address` predefined variable to dynamically assign a secondary IP address to the unnumbered interface. You cannot configure a static preferred source address.

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

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

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

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

PE router to protect the best path. When BFD is enabled on the BGP session between the CE and the primary PE router, with local traffic flowing from another CE connected with the primary PE to this CE, after bringing the interface down on the best path, the local repair will be triggered by BFD session down, but it might fail due to a timing issue. This will cause slow converge and unexpected traffic drop. [PR1098961](#)

- When the BFD is running on multi LU (lookup chip) Packet Forwarding Engine (such as MPC3 or MPC4), incoming BFD packet might be processed with a firewall filter on different logical-routers's loopback interface. If the firewall filter is discarding/rejecting BFD, the packets will be dropped incorrectly. [PR1099608](#)
- On MX Series-based platform, before creating a new unicast next hop, there is a check to see if there is at least 512k DoubleWords (DW) free. So, even the attempting NH requires only a small amount of memory (for example, < 100 DWs), if there is no such enough free DWs (that is, 512k), the check will fail and the end result is that the control plane will quit adding this NH prematurely - stopping at ~80% of capacity. With the fix, it will check for 64k free DWs which is lower reference watermark for available resource, thereby ensuring that can allocate resource. [PR1099753](#)
- From Junos OS Release 14.1 and above, IPv6 mobility packets with Heartbeat option that the length of the mobility header (including the ethernet encapsulation and main IPv6 header) extends beyond 128 Bytes will be discarded as bad IPv6 option packet due to a logic error in packet handling. [PR1100442](#)
- On MX Series platform, when using the 64-bit image, if the configuration statement "source-address" is configured for the "radius-server" as the following, the RADIUS request may not be sent to RADIUS server due to the failure of setting the "source-address" on the device. `user@re0> show configuration system radius-server .. source-address 10.1.1.1; <<<<< The configuration statement that may cause the issue` [PR1103517](#)
- On T4000 platform with FPC Type-5 equipped, after performing unified ISSU, due to the fact that only 6 out of 16 temperature sensors may get initialized, the temperature reading for the line card may be shown as "Absent". [PR1104240](#)
- Due to a software defect found in 13.3R7.3 and 14.1R5.4, Juniper Networks strongly discourage the use of Junos OS software version 13.3R7.3 on routers with MQ-based MPC. This includes MX Series with MPC1, MPC2; all mid-range MX Series. [PR1108826](#)

Routing Policy and Firewall Filters

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

Routing Protocols

- Support for the Pragmatic General Multicast protocol (daemon pgmd) is being phased out from Junos OS. In Junos OS Release 14.2, the CLI is now hidden (although the component is still there and configurable). In Junos OS Release 15.1 the code and its corresponding CLI are removed. [PR936723](#)
- In PIM multicast-only fast reroute (MoFRR) environment, when issuing CLI command "show multicast route extensive" on egress edge router, due to missing null check while showing label information for reverse-path forwarding (RPF) nexthop, an error might be seen in the output of the command. In addition, the routing protocol process (rpd) may crash on the device. [PR983140](#).
- For the pim nbr which is not directly connected (that is, nbr on unnumbered interface, or p2p interface with different subnet), pim join is not able to find the correct upstream nbr which results in join not propagating to the upstream nbr . show command for pim join shows upstream nbr "unknown" . Issue is present in the 15.1R1 release. [PR1069896](#)
- In mutli-topologies IS-IS scenario, there is huge difference between estimated free bytes and actual free bytes when generating LSP with IPv6 Prefix. It might cause LSP fragment exhaustion. [PR1074891](#).
- There are two issues in the PR: (1) In multicast environment, Incoming interface list (IIF) list has only RPF interface, designated forwarder (DF) winners are not added in the list in backup Routing Engine. (2) "Number of downstream interfaces" in show pim join extensive is not accounting Pseudo-VXLAN interface. [PR1082362](#)
- If a policy statement referred to a routing-table, but the corresponding routing instance is not fully configured (ie. no instance-type), commit such configuration might cause the rpd process to crash. [PR1083257](#).
- With Multicast Source Discovery Protocol (MSDP) and nonstop active routing (NSR) configured on the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP), the rpd process might permanently get stuck when multicast traffic received shortly after Routing Engines switchover. [PR1083385](#)
- When there are a number of secondary BGP routes in inet.0, an SNMP walk of inet.0 by the bgp4 MIB can cause a core if the corresponding primary routes are being deleted. [PR1083988](#).
- 1. configure the ospf and ospf3 in all routers 2. configure node protection 3. check for 22.1.1.0 any backup is present 4. enable pplfa all 5. check for 22.1.1.0 any pplfa backup is present through r2 we are not seeing any pplfa backup for 22.1.1.0 [PR1085029](#)
- When BGP route is leaked to a routing-instance and there is an import policy to overwrite the route preference, if damping is also configured in BGP, the BGP routes which were copied to second table cannot be deleted after routes were deleted in master table. This is a day-1 issue. [PR1090760](#)
- When removing BGP Prefix-Independent Convergence (PIC) from the configuration, the expected behavior is that any protected path would become unprotected. But in this case, the multipath entry that contains the protection path (which is supposed to be removed) remains active, until BGP session flaps or the route itself flaps. As a workaround, we can use "commit full" command to correct or to commit. [PR1092049](#)

