

Release Notes: Junos[®] OS Release 15.1R6 for the ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series

4 October 2018

Contents

Introduction	8
Junos OS Release Notes for ACX Series	8
New and Changed Features	8
Hardware	9
Class of Service	9
Firewall Filters	10
Interfaces and Chassis	11
Installation	17
Layer 2 Features	17
Management	22
Routing	23
Security	23
Subscriber Access Management	23
Timing and Synchronization	23
Changes in Default Behavior and Syntax	25
General Routing	25
Interfaces and Chassis	25
Management	25
Known Behavior	26
Known Issues	26
Class of Service	27
Firewall Filters	28
Interfaces and Chassis	29
Integrated Routing and Bridging	31
Layer 2 Services	33
MPLS Applications	34

Network Management	34
Statistics	34
Timing and Synchronization	34
Resolved Issues	35
Resolved Issues	35
Documentation Updates	35
Migration, Upgrade, and Downgrade Instructions	36
Upgrade and Downgrade Support Policy for Junos OS Releases	36
Product Compatibility	37
Hardware Compatibility	37
Junos OS Release Notes for EX Series Switches	37
New and Changed Features	37
Hardware	38
Authentication and Access Control	39
Interfaces and Chassis	40
Junos OS XML API and Scripting	42
Management	42
MPLS	43
Network Management and Monitoring	43
Port Security	43
Software Installation and Upgrade	44
Spanning-Tree Protocols	44
Changes in Behavior and Syntax	45
Dynamic Host Configuration Protocol	45
Management	45
Virtual Chassis	45
Known Behavior	46
Authentication and Access Control	46
Infrastructure	47
Interfaces and Chassis	47
J-Web	48
Layer 2 Features	50
MPLS	50
Multicast Protocols	50
Network Management and Monitoring	51
Platform and Infrastructure	51
Port Security	51
Routing Protocols	52
Software Installation and Upgrade	52
User Interface and Configuration	52
Virtual Chassis	53
Known Issues	53
Authentication and Access Control	53
General routing	54
Layer 2 Features	54
Security	54
Resolved Issues	54
Resolved Issues: Release 15.1R6	55
Resolved Issues: Release 15.1R5	60

Resolved Issues: Release 15.1R4	64
Resolved Issues: Release 15.1R3	67
Resolved Issues: Release 15.1R2	73
Resolved Issues: Release 15.1R1	76
Documentation Updates	76
Changes to the Junos OS for EX Series Documentation	77
Errata in the Junos OS for EX Series Documentation	77
Migration, Upgrade, and Downgrade Instructions	77
Upgrade and Downgrade Support Policy for Junos OS Releases	77
Product Compatibility	78
Hardware Compatibility	78
Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, and T Series Core Routers	80
New and Changed Features	80
Hardware	81
Bridging and Learning	81
Class of Service (CoS)	82
High Availability (HA) and Resiliency	83
Interfaces and Chassis	86
IPv6	90
Junos OS XML API and Scripting	90
Layer 2 Features	91
Management	93
MPLS	94
Multicast	96
Network Management and Monitoring	97
Routing Policy and Firewall Filters	99
Routing Protocols	100
Services Applications	102
Software Defined Networking	108
Software Installation and Upgrade	109
Software Licensing	109
Subscriber Management and Services (MX Series)	112
System Logging	129
User Interface and Configuration	129
VPNs	129
Changes in Behavior and Syntax	131
Authentication, Authorization and Accounting	132
Class of Service (CoS)	132
General Routing	133
High Availability (HA) and Resiliency	133
IPv6	135
Junos OS XML API and Scripting	135
Layer 2 Features	136
Layer 2 VPNs	136
Management	136
MPLS	136
Multicast	137
Network Management and Monitoring	137

Platform and Infrastructure	138
Routing Policy and Firewall Filters	138
Routing Protocols	139
Security	142
Services Applications	144
Subscriber Management and Services (MX Series)	146
System Logging	157
System Management	164
User Interface and Configuration	164
Virtual Chassis	165
VLAN Infrastructure	165
VPNs	165
Known Behavior	165
Hardware	166
Forwarding and Sampling	167
Interfaces and Chassis	167
MPLS	167
Network Management and Monitoring	167
Subscriber Management and Services (MX Series)	167
System Logging	169
VPNs	170
Known Issues	170
Forwarding and Sampling	171
General Routing	172
Infrastructure	175
Interfaces and Chassis	175
J-Web	175
Layer 2 Ethernet Services	176
MPLS	176
Network Management and Monitoring	177
Platform and Infrastructure	177
Routing Protocols	180
Services Applications	181
Subscriber Access Management	181
User Interface and Configuration	181
VPNs	182
Infrastructure	182
Interfaces and Chassis	183
Layer 2 Ethernet Services	184
Multiprotocol Label Switching (MPLS)	184
Network Management and Monitoring	185
Platform and Infrastructure	186
Routing Protocols	188
Services Applications	190
Subscriber Access Management	191
User Interface and Configuration	191
VPNs	191

Resolved Issues	192
Resolved Issues: 15.1R6	192
Resolved Issues: 15.1R5	212
Resolved Issues: 15.1R4	234
Resolved Issues: 15.1R3	253
Resolved Issues: 15.1R2	289
Documentation Updates	312
Adaptive Services Interfaces Feature Guide for Routing Devices	312
Advanced Subscriber Management Provisioning Guide	313
Broadband Subscriber Sessions Feature Guide	313
Broadband Subscriber VLANs and Interfaces Feature Guide	314
High Availability Feature Guide	314
IPv6 Neighbor Discovery Feature Guide for Routing Devices	315
Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices	315
MPLS Applications Feature Guide for Routing Devices	316
Overview for Routing Devices	317
Release Notes	317
Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices	318
Security Services Administration Guide for Routing Devices	318
Standards Reference	318
Subscriber Management Access Network Guide	318
Subscriber Management Provisioning Guide	319
Tunnel and Encryption Services Interfaces	319
User Access and Authentication Guide for Routing Devices	319
VPNs Library for Routing Devices	319
Migration, Upgrade, and Downgrade Instructions	320
Basic Procedure for Upgrading to Release 15.1	321
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x)	322
Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)	323
Upgrade and Downgrade Support Policy for Junos OS Releases	325
Upgrading a Router with Redundant Routing Engines	325
Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1	326
Upgrading the Software for a Routing Matrix	327
Upgrading Using Unified ISSU	328
Downgrading from Release 15.1	329
Product Compatibility	329
Hardware Compatibility	329
Junos OS Release Notes for PTX Series Packet Transport Routers	331
New and Changed Features	331
High Availability and Resiliency (HA)	332
Interfaces and Chassis	332
IPv6	333
Junos OS XML API and Scripting	333
Management	334
MPLS	335

Routing Protocols	335
Software Licensing	336
User Interface and Configuration	339
VPNs	339
Changes in Behavior and Syntax	340
High Availability (HA) and Resiliency	340
IPv6	341
Junos OS XML API and Scripting	341
Management	341
Network Management and Monitoring	341
Routing Policy and Firewall Filters	342
Routing Protocols	342
User Interface and Configuration	342
Known Behavior	343
System Logging	343
Known Issues	343
General Routing	344
Infrastructure	345
Interfaces and Chassis	345
MPLS	345
Platform and Infrastructure	345
Routing Protocols	345
Resolved Issues	346
Resolved Issues: 15.1R6	346
Resolved Issues: 15.1R5	348
Resolved Issues: 15.1R4	351
Resolved Issues: 15.1R3	354
Resolved Issues: 15.1R2	357
Documentation Updates	360
High Availability Feature Guide	360
IPv6 Neighbor Discovery Feature Guide	360
Migration, Upgrade, and Downgrade Instructions	361
Upgrading Using Unified ISSU	361
Upgrading a Router with Redundant Routing Engines	361
Basic Procedure for Upgrading to Release 15.1	361
Product Compatibility	365
Hardware Compatibility	365
Junos OS Release Notes for the QFX Series	366
New and Changed Features	366
Management	366
Network Management and Monitoring	368
Spanning-Tree Protocols	368
User Interface and Configuration	368
Changes in Behavior and Syntax	369
Interfaces and Chassis	369
Routing Protocols	370
Known Behavior	370
High Availability (HA) and Resiliency	371
Interfaces and Chassis	371

Layer 2 Features	371
Multicast Protocols	372
Multiprotocol Label Switching (MPLS)	372
Platform and Infrastructure	372
Routing Policy	372
Routing Protocols	372
Software-Defined Networks (SDN)	373
Software Installation and Upgrade	373
Spanning-Tree Protocols	373
Virtual Chassis and Virtual Chassis Fabric (VCF)	373
VPNs	373
Known Issues	374
Routing Protocols	374
Security	374
Spanning Tree Protocols	374
Virtual Chassis and Virtual Chassis Fabric	374
Resolved Issues	375
Resolved Issues: Release 15.1R6	375
Resolved Issues: Release 15.1R5	377
Resolved Issues: Release 15.1R4	380
Resolved Issues: Release 15.1R3	383
Documentation Updates	390
Migration, Upgrade, and Downgrade Instructions	390
Upgrading Software on QFX3500, QFX3600, and QFX5100 Standalone Switches	390
Performing an In-Service Software Upgrade (ISSU) on the QFX5100 Switch	392
Product Compatibility	394
Hardware Compatibility	394
Third-Party Components	396
Upgrading Using Unified ISSU	396
Compliance Advisor	396
Finding More Information	396
Documentation Feedback	396
Requesting Technical Support	398
Self-Help Online Tools and Resources	398
Opening a Case with JTAC	398
Revision History	399

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 15.1R6 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 15.1R6 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 https://www.juniper.net/documentation/product/en_US/junos-os.

- [New and Changed Features on page 8](#)
- [Changes in Default Behavior and Syntax on page 25](#)
- [Known Behavior on page 26](#)
- [Known Issues on page 26](#)
- [Resolved Issues on page 35](#)
- [Documentation Updates on page 35](#)
- [Migration, Upgrade, and Downgrade Instructions on page 36](#)
- [Product Compatibility on page 37](#)

New and Changed Features

This section describes the features and enhancements in Junos OS Release 15.1R6 for ACX Series Universal Metro Routers.

- [Hardware on page 9](#)
- [Class of Service on page 9](#)
- [Firewall Filters on page 10](#)
- [Interfaces and Chassis on page 11](#)
- [Installation on page 17](#)
- [Layer 2 Features on page 17](#)
- [Management on page 22](#)
- [Routing on page 23](#)
- [Security on page 23](#)
- [Subscriber Access Management on page 23](#)
- [Timing and Synchronization on page 23](#)

Hardware

- **ACX Series Universal Metro Router**—Starting with Junos OS Release 15.1R3, Junos OS supports the following Juniper Networks ACX Series Universal Metro Routers:
 - [ACX1000 and ACX1100 Universal Metro Router](#)
 - [ACX2000 and ACX2100 Universal Metro Router](#)
 - [ACX2200 Universal Metro Router](#)
 - [ACX4000 Universal Metro 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 Metro 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 Metro 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
<code>first-fragment</code>	Matches if packet is the first fragment
<code>is-fragment</code>	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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro Routers supports Layer 2 circuit and Layer 2 VPN Ethernet virtual connection (EVC) types.

Junos OS for ACX Series Universal Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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 Metro 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.

- See Also**
- [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Behavior on page 26](#)
 - [Known Issues on page 26](#)
 - [Resolved Issues on page 35](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 36](#)
 - [Product Compatibility on page 37](#)

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 15.1R6 for the ACX Series Universal Metro Routers.

- [General Routing on page 25](#)
- [Interfaces and Chassis on page 25](#)
- [Management on page 25](#)

General Routing

- Routing commands in 64-bit mode enabled by default—Starting in Junos OS Release 15.1F3 and later, routing commands in 64-bit mode are enabled by default on systems that support that mode and that have at least 16 GB of RAM.

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.

Management

- **Support for status deprecated statement in YANG modules (ACX Series)**—Starting with Junos OS Release 15.1R5, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

See Also • [New and Changed Features on page 8](#)

- [Known Behavior on page 26](#)
- [Known Issues on page 26](#)
- [Resolved Issues on page 35](#)
- [Documentation Updates on page 35](#)
- [Migration, Upgrade, and Downgrade Instructions on page 36](#)
- [Product Compatibility on page 37](#)

Known Behavior

There are no known limitations in Junos OS Release 15.1R6 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.

- See Also**
- [New and Changed Features on page 8](#)
 - [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Issues on page 26](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 35](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 36](#)
 - [Product Compatibility on page 37](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R6 for the ACX Series Universal Metro 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 27](#)
- [Firewall Filters on page 28](#)
- [Interfaces and Chassis on page 29](#)
- [Integrated Routing and Bridging on page 31](#)
- [Layer 2 Services on page 33](#)
- [MPLS Applications on page 34](#)
- [Network Management on page 34](#)
- [Statistics on page 34](#)
- [Timing and Synchronization on page 34](#)

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 Metro 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 Metro 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 Metro 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 Metro 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](#)

- See Also**
- [New and Changed Features on page 8](#)
 - [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Behavior on page 26](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 35](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 36](#)
 - [Product Compatibility on page 37](#)

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 35](#)

Resolved Issues

There are no resolved issues in 15.1R6.

- See Also**
- [New and Changed Features on page 8](#)
 - [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Behavior on page 26](#)
 - [Known Issues on page 26](#)
 - [Documentation Updates on page 35](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 36](#)
 - [Product Compatibility on page 37](#)

Documentation Updates

There are no errata or changes in Junos OS Release 15.1R6 for the ACX documentation.

- See Also**
- [New and Changed Features on page 8](#)
 - [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Behavior on page 26](#)
 - [Known Issues on page 26](#)

- [Resolved Issues on page 35](#)
- [Migration, Upgrade, and Downgrade Instructions on page 36](#)
- [Product Compatibility on page 37](#)

Migration, Upgrade, and Downgrade Instructions

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Universal Metro 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 36](#)

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

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

- See Also**
- [New and Changed Features on page 8](#)
 - [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Behavior on page 26](#)
 - [Known Issues on page 26](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 35](#)
 - [Product Compatibility on page 37](#)

Product Compatibility

- [Hardware Compatibility on page 37](#)

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 <https://pathfinder.juniper.net/feature-explorer/>.

- See Also**
- [New and Changed Features on page 8](#)
 - [Changes in Default Behavior and Syntax on page 25](#)
 - [Known Behavior on page 26](#)
 - [Known Issues on page 26](#)
 - [Resolved Issues on page 35](#)
 - [Documentation Updates on page 35](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 36](#)

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 15.1R6 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 https://www.juniper.net/documentation/product/en_US/junos-os.

- [New and Changed Features on page 37](#)
- [Changes in Behavior and Syntax on page 45](#)
- [Known Behavior on page 46](#)
- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 77](#)
- [Product Compatibility on page 78](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R6 for the EX Series.



NOTE: The following EX Series platforms are supported in Junos OS Release 15.1R6: 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 38](#)
 - [Authentication and Access Control on page 39](#)
 - [Interfaces and Chassis on page 40](#)
 - [Junos OS XML API and Scripting on page 42](#)
 - [Management on page 42](#)
 - [MPLS on page 43](#)
 - [Network Management and Monitoring on page 43](#)
 - [Port Security on page 43](#)
 - [Software Installation and Upgrade on page 44](#)
 - [Spanning-Tree Protocols on page 44](#)

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 fallback 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 an 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

- **Half-duplex link support (EX4300 switches)**—Starting with Junos OS 15.1R4, half-duplex communication is supported on all built-in network copper ports on EX4300 switches. *Full-duplex* communication means that both ends of the communication can send and receive signals at the same time. *Half-duplex* is also bidirectional communication, but signals can flow in only one direction at a time. Half-duplex is configured by default on EX4300 switches. If the link partner is set to autonegotiate the link, then the link is autonegotiated to full duplex or half duplex. If the link is not set

to autonegotiation, then the EX4300 link defaults to half-duplex unless the interface is explicitly configured for full duplex.

To explicitly configure full duplex:

```
[edit]
user@switch# set interfaces interface-name speed 10m-or-100m
[edit]
user@switch# set interfaces interface-name ether-options no-auto-negotiate
```

To verify a half-duplex setting, issue *one* of:

```
user@switch> show interfaces interface-name media
user@switch> show interfaces interface-name extensive
```

To query the OID, issue:

```
user@switch> show snmp mib get dot3StatsDuplexStatus.SNMP-ifIndex
```

[See “[Documentation Updates](#)” on page 76.]

- **LACP minimum link support on LAGs (EX9200 switches)**—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](#).]

Network Management and Monitoring

- **MIB support for media attachment unit (MAU) information (EX2200, EX3300)**—Starting with Junos OS Release 15.1R4, EX2200 and EX3300 switches support standard and enterprise-specific MIBs that allow users to gather information about MAUs connected to those switches. The switches populate the entityMIB (RFC 4133) and entityStateMIB (RFC 4268) standard SNMP MIBs, and a new MIB table, ifJnxMediaTable, which is part of the Juniper enterprise-specific Interface MIB extensions. The objects in ifJnxMediaTable represent MAU information such as media type, connector type, link mode, and link speed. Users can gather this information using the Junos OS CLI command **show snmp mib** or other remote SNMP MIB object access methods.

[See [SNMP MIB Explorer](#).]

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 shut down 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 switches, 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, ELS 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 ELS 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](#).]

- See Also**
- [Changes in Behavior and Syntax on page 45](#)
 - [Known Behavior on page 46](#)
 - [Known Issues on page 53](#)
 - [Resolved Issues on page 54](#)
 - [Documentation Updates on page 76](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 77](#)
 - [Product Compatibility on page 78](#)

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 15.1R6 for the EX Series.

- [Dynamic Host Configuration Protocol on page 45](#)
- [Management on page 45](#)
- [Virtual Chassis on page 45](#)

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.

Management

- **Support for status deprecated statement in YANG modules (EX Series)**—Starting with Junos OS Release 15.1R5, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

Virtual Chassis

- **Increased time to rejoin Virtual Chassis after a member is rebooted (EX2200, EX3300, EX4200, EX6200, and EX8200 Virtual Chassis)**—Starting in Junos OS Release 15.1R3, when one or more member switches in an EX2200, EX3300, EX4200, EX6200, or EX8200 Virtual Chassis are rebooted, the Virtual Chassis master's delay time before reinstating the rebooted switch as a member in the Virtual Chassis is increased from two minutes to ten minutes. As a result, after rebooting a Virtual Chassis member, up to 15 or 20 minutes total elapsed time might be required for the member to completely rejoin the Virtual Chassis. The increased delay time allows the Virtual Chassis to correctly rebuild its Virtual Chassis port (VCP) adjacency information, and avoids unexpected mastership election contention or failure of the Virtual Chassis to re-form.

- See Also**
- [New and Changed Features on page 37](#)
 - [Known Behavior on page 46](#)
 - [Known Issues on page 53](#)
 - [Resolved Issues on page 54](#)
 - [Documentation Updates on page 76](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 77](#)
 - [Product Compatibility on page 78](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R6 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](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [J-Web](#)
- [Layer 2 Features](#)
- [MPLS](#)
- [Multicast Protocols](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Port Security](#)
- [Routing Protocols](#)
- [Software Installation and Upgrade](#)
- [User Interface and Configuration](#)
- [Virtual Chassis](#)

Authentication and Access Control

- On EX4300 switches, a maximum of 5K supplicants is supported for dot1xd. [PR962292](#)
- 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](#)
- On an EX4300 or a QFX5100 switch, a MAC address that is specified as part of a MAC-based VLAN is authenticated on an interface, for example, xe-1/1/1, on which 802.1X authentication in multiple supplicant mode is configured. However, the same

MAC address might not be authenticated on another interface, for example, xe-2/1/1, if the MAC address moves to interface xe-2/1/1 from interface xe-1/1/1. [PR1007589](#)

Infrastructure

- On EX Series switches, ARP reply packets might get dropped when the switch receives reverse-path forwarding (RPF) multicast failure packets at a high rate (for example, 300 pps). As a workaround, create a static ARP entry for the next-hop device. [PR1007438](#)
- System logging (syslog) messages for EX9200-MPC line cards include error messages on FPC initialization. Initialization can be triggered by FPC restart, insertion and removal, or power off and on. The message is **Error "kernel: GENCFG: op 32 (Resync blob) failed; err 7 (Doesn't Exist)"**. This has no functional impact. [PR1171487](#)
- On EX3300 and EX4200 switches, DHCPv6 packets are duplicated with option18 configured (one packet with option 18 and one without option 18) when switches are configured with **dhcpv6-option18 use-option82**. This is an expected behavior. [PR1184593](#)

Interfaces and Chassis

- On an EX2200 Virtual Chassis with three members, if you configure nine link aggregation groups (LAGs) and eight interfaces per LAG bundle, the LACP links might move down and up continuously. As a workaround, configure eight link aggregation groups and eight interfaces per LAG bundle. [PR1030809](#)
- On EX9200 switches configured with an MC-LAG, the Inter-Chassis Control Protocol (ICCP) might flap if you configure another interchassis link (ICL) that is on new multichassis aggregated Ethernet (MC-AE) interfaces. [PR1046022](#)
- On EX9200 switches on which a MAC limit is configured with **packet-action log**, a packet drop might occur when **interface-mac-limit** is configured with **mac-table-size** on a specific VLAN or on a global VLAN hierarchy. [PR1076546](#)
- On EX9200 switches, if you configure **mac-move-limit** with **packet-action shutdown** on a VLAN that includes an MC-AE interface and an access interface, the packet action is not performed if traffic hits the limit between the MC-AE interface and the access interface. [PR1079383](#)
- On EX9200 switches, if you configure **mac-move-limit** with **packet-action shutdown** on a VLAN that includes two members of a multichassis link aggregation group (MC-LAG) AE interface, if traffic hits the limit between the two MC-AE interfaces, a peer link belonging to one of the MC-AE interfaces might go down and only 50 percent of the traffic might reach its destination. [PR1079436](#)
- On EX9200 switches, unified in-service software upgrade (ISSU) might not work properly for an upgrade to Junos OS Release 15.1. As a workaround, manually upgrade the Routing Engine. [PR1091610](#)

- 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 Inter-Chassis Control Protocol (ICCP) goes down and comes back up. [PR1107001](#)
- If an Inter-Chassis Control Protocol (ICCP) interface on an EX9200 switch in an MC-LAG Active-Active topology is disabled and then reenabled, traffic could be dropped for more than 2 seconds. [PR1173923](#)

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](#)

- In the J-Web interface, in the Port Security Configuration page, configuring the **action** option when you configure the **MAC limit** option is mandatory, even though configuring an action value is not mandatory in the CLI. [PR434836](#)
- On EX4200 switches, in the J-Web interface, if you try to change the position of columns using the drag-and-drop method, only the column headers move to the new position instead of the entire column in the OSPF Global Settings table in the OSPF Configuration page, the Global Information table in the BGP Configuration page, and the Add Interface window in the LACP (Link Aggregation Control Protocol) Configuration page. [PR465030](#)
- When a large number of static routes are configured and you have navigated to pages other than page 1 in the Route Information table in the Static Routing monitoring page in the J-Web interface (Monitor > Routing > Route Information), changing the Route Table to query other routes refreshes the page, but does not return to page 1. For example, if you run a query from page 3 and the new query returns very few results, the Results table continues to display page 3 and shows no results. To view the results, navigate to page 1 manually. [PR476338](#)
- In the J-Web interface for EX4500 switches, the Port Configuration page (Configure > Interfaces > Ports), the Port Security Configuration page (Configure > Security > Port Security), and the Filters Configuration page (Configure > Security > Filters) display features that are not supported on EX4500 switches. [PR525671](#)
- When you open a J-Web interface session using HTTPS, enter a username and a password, and then click the Login button, the J-Web interface takes 20 seconds longer to launch and load the Dashboard page than it does if you use HTTP. [PR549934](#)
- If you access the J-Web interface by using an HTTPS connection through the Microsoft Internet Explorer Web browser, you might not be able to download and save reports from some pages on the Monitor, Maintain, and Troubleshoot tabs. Some affected pages are at these locations:

- Maintain > Files > Log Files > Download
- Maintain > Config Management > History
- Maintain > Customer Support > Support Information > Generate Report
- Troubleshoot > Troubleshoot Port > Generate Report
- Monitor > Events and Alarms > View Events > Generate Report
- Monitor > Routing > Route Information > Generate Report

As a workaround, use the Mozilla Firefox Web browser to download and save reports while using an HTTPS connection. [PR566581](#)

- 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](#)
- On an EX4300 Virtual Chassis, if you renumber the Virtual Chassis members while there is an active J-Web session, a socket error might be created. As a workaround, refresh the J-Web session. [PR857269](#)
- On EX Series switches, the subscriber management infrastructure daemon (smid) might randomly crash when the smid daemon is interleaved with another daemon that is attempting to access the same shared memory. [PR1082211](#)
- On an EX4600 Virtual Chassis, if lossless traffic is passing through a switch in the linecard role over a 10-gigabit SFP+ link configured as a Virtual Chassis port (VCP), traffic on the link might be dropped when the link is congested. [PR1006974](#)
- On EX Series switches except EX4600, if you configure an IPv4 GRE interface on an IPv6 interface, the GRE tunnel might not work properly. Traffic is not forwarded through the tunnel. [PR1008157](#)
- The J-Web dashboard might take longer than usual to load depending on the number of EX8200 Virtual Chassis members, due to time taken for collecting CLI responses. [PR806803](#)

Layer 2 Features

- On EX Series switches, after a switch reboot, a Q-in-Q tunneling interface might not function as expected. The problem occurs when the interface is a member of a PVLAN with mapping set to swap and is also a member of a non-private VLAN. The PVID of the access interface does not get set when the PVLAN is configured before the non-private VLAN. The problem does not occur when the non-private VLAN is configured before the PVLAN. [PR937927](#)
- On ELS (Enhanced Layer 2 Software) platforms (including EX4300, EX4600, EX9200, QFX3500, QFX3600, and QFX5100), if Q-in-Q tunneling is enabled, if you configure an RTG (redundant trunk group) on a Q-in-Q interface, the RTG configuration cannot be applied; there is a commit check error. [PR1134126](#)
- On EX4500 Virtual Chassis, when one member of the Virtual Chassis is switched off, ERPS should be reinitialized on the other members. However, because the interfaces are on the member ERPS ring that is not active anymore, ERPS cannot complete initialization properly and it stays in the init state. Thus, the rest of the interfaces do not converge to a proper state. This is expected behavior. If there is a requirement to have complete ERPS support in a ring topology and to perform a mastership failover test on a Virtual Chassis with ERPS, then Interfaces of the Virtual Chassis that are part of the ERPS link should be configured as aggregated Ethernet (AE) interfaces. Ideally physical interfaces that are part of this AE interface should be spread across all members of the Virtual Chassis. However, this is not necessary—the AE interface can contain only one physical interface and the mastership failover will still work properly. [PR1235062](#)

MPLS

- On EX4600 switches, user-to-network (UNI) interfaces that have over 100 pseudowires might not function correctly. Up to 100 pseudowires are supported in active/backup configurations (cold standby). If more than 100 active and backup pseudowires are configured, traffic might not be forwarded correctly after a provider edge (PE) switch is either rebooted or disabled then reenabled. [PR1048500](#)

Multicast Protocols

- On EX4550 switches, if you configure IGMP on all interfaces and create a large number of multicast groups, the maximum scale for IGMP can be achieved on some interfaces but not all interfaces. [PR1025169](#)
- On EX9200 switches, multicast traffic might fail when the source is on an ordinary VLAN and the receiver is on a PVLAN with a primary VLAN ID, with both source and receiver on the same switch. [PR1028869](#)
- On Virtual Chassis models EX2200, EX3300, EX4200, EX4500, EX4550, and EX8200, Layer 3 multicast traffic does not flow if VLAN pruning is enabled for the upstream interface and the VLAN does not have a member on the device on which the downstream interface resides. As a workaround, disable VLAN pruning for the upstream interface if the device where the downstream interface resides does not have a member for that VLAN. [PR1156014](#)

Network Management and Monitoring

- This is a limitation with the physical layer (being used in EX4550-32F), while reading the SFP-T optics registers (16-bit register), hence there is a delay while doing an SNMP walk or an SNMP GET request for interface-specific MIBs. [PR832071](#)
- 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](#)
- On EX8200 switches, some sFlow data might have incorrect input and output interface index values. [PR1051435](#)

Platform and Infrastructure

- You cannot connect EX2200-C-12P-2G switches to the prestandard Cisco IP Phone 7960 using a straight cable. As a workaround, use a crossover cable. [PR726929](#)
- On EX4300 switches, Ethernet ring protection (ERP) fails if the control VLAN is replaced with a different VLAN at runtime. [PR817456](#)
- On EX4300 switches, despite an administrative link being down, child members of an aggregated Ethernet group that is part of a multicast downstream IRB VLAN might be programmed into a multicast route index in the Packet Forwarding Engine, resulting in the failure of multicast replication of packets for some VLANs. [PR880769](#)
- On EX4300 switches, if multicast data packets that fail an RPF check are received on a nonshared tree, the packets might be trapped on the Routing Engine at a high rate, resulting in poor PIM convergence. [PR911649](#)
- On EX4300 switches, in an egress router-based firewall filter, IPv6 Layer 4 headers (of ICMP type) might not work. [PR912483](#)
- On EX9200 Virtual Chassis, commit errors might occur if commits are done frequently. [PR1188816](#)

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 EX4600 switches, when LACP is configured together with MACsec, the links in the bundle might not all work. Rebooting the switch might solve the problematic links but might also create the same issue on other child interfaces. [PR1093295](#)
- On EX4300 switches, when **storm-control** or **storm-control-profiles** with **action-shutdown** is configured, if the storm-triggered traffic is control traffic such as LACP, the physical interface might be put into an STP blocking state rather than turned down, so valid control traffic might be trapped to the control plane and unrelated interfaces might be set down as an LACP timeout. [PR1130099](#)

Routing Protocols

- On EX4300, EX4600, and QFX Series switches, a Bidirectional Forwarding Detection (BFD) session might not come up when BFD version 0 is configured. As a workaround, deactivate or delete the version configuration. [PR1076052](#)

Software Installation and Upgrade

- On EX Series or QFX Series Virtual Chassis or Virtual Chassis Fabric (VCF), nonstop software upgrade (NSSU) cannot be used to upgrade from a Junos OS Release 14.1X53 image to a Junos OS Release 15.1 or later image. [PR1087893](#)
- On EX4600, QFX3500, and QFX5100 switches, the amount of time that it takes for Zero Touch Provisioning to complete might be lengthy because TFTP might take a long time to fetch required data. [PR980530](#)
- 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](#)
- Substantial traffic losses might occur when executing a nonstop software upgrade (NSSU) on a mixed EX4200 and EX4500 Virtual Chassis or on an EX3300 Virtual Chassis, an EX6200 switch, an EX8200 switch, or an EX8200 Virtual Chassis. [PR1062960](#)
- 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](#)
- On EX4300 switches, traffic might be lost for Layer 3 protocols (such as RIP, OSPF, BGP, and VRRP) during a nonstop system upgrade (NSSU). [PR1065405](#)
- 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](#)
- 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](#)

User Interface and Configuration

- On EX8200 Virtual Chassis, if you are using the Virtual Chassis wizard in the J-Web interface in the Mozilla Firefox version 3.x browser and select more than six port pairs from the same member to convert from VCPs to network ports, the wizard might display incorrect port conversion status. Also, if you double-click **Next** after deleting an active member in the Members page, the J-Web interface might stop responding. [PR796584](#)

- If you uninstall the J-Web Platform package by using the CLI, reinstalling the Application package does not restore J-Web. [PR1026308](#)

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](#)
- Automatic software update is not supported on an EX2200 Virtual Chassis.

- See Also**
- [New and Changed Features on page 37](#)
 - [Changes in Behavior and Syntax on page 45](#)
 - [Known Issues on page 53](#)
 - [Resolved Issues on page 54](#)
 - [Documentation Updates on page 76](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 77](#)
 - [Product Compatibility on page 78](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R6 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 on page 53](#)
- [General routing on page 54](#)
- [Layer 2 Features on page 54](#)
- [Security on page 54](#)

Authentication and Access Control

- On an EX4300 Virtual Chassis that is configured for 802.1X authentication, an invalid supplicant might remain in a connecting state instead of moving to a held state. [PR1149008](#)
- On a dot1x-enabled interface, sometimes when you log in, log off, and then log in within a short interval (within subseconds), the logical interface plus the bridge domain or VLAN remain in a pending state, and you will not be able to access the network. As a workaround, restart the l2-learning process to recover the port/interface from the problematic state. [PR1230073](#)

General routing

- On EX Series switches that run Enhanced Layer 2 Software (ELS), when an interface is removed from a private VLAN (PVLAN) and then added back, the corresponding MAC entry might not be deleted from the Ethernet table. [PR1036265](#)

Layer 2 Features

- On an EX9200-6QS line card, storm control might not work for multicast traffic. [PR1191611](#)
- The dest-MAC validation feature uses MLP handshakes to detect stale destination-MAC addresses. If a stale MAC address is detected, the system automatically deletes it. The deletion of destination MACs does not cause traffic drops, as the next packet is flooded and valid MACs are relearned. On an EX9200 Virtual Chassis, MLP handshakes are occasionally dropped across Virtual Chassis members. This drop is random and occurs only when a source-MAC and its related dest-MACs are on different member chassis. This causes intermittent dest-MAC deletion and flooding; however, no packet drop results because of this. [PR1249788](#)

Security

- On EX4300, EX4600, and QFX5100 switches, when a VLAN is mirrored, the mirrored packets might contain 38 additional bytes. The IP address in this packet is randomly generated and might appear as one of many existing, valid IP addresses on the Internet. It might appear as ERSPAN as well, which is a proprietary non-Juniper protocol. These addresses and packet types can be ignored. They might appear as alerts in certain IDPs or IDSs and in packet analyzer applications, which you can ignore. [PR1170589](#)

- See Also**
- [New and Changed Features on page 37](#)
 - [Changes in Behavior and Syntax on page 45](#)
 - [Known Behavior on page 46](#)
 - [Resolved Issues on page 54](#)
 - [Documentation Updates on page 76](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 77](#)
 - [Product Compatibility on page 78](#)

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 15.1R6 on page 55](#)
- [Resolved Issues: Release 15.1R5 on page 60](#)
- [Resolved Issues: Release 15.1R4 on page 64](#)
- [Resolved Issues: Release 15.1R3 on page 67](#)
- [Resolved Issues: Release 15.1R2 on page 73](#)
- [Resolved Issues: Release 15.1R1 on page 76](#)

[Resolved Issues: Release 15.1R6](#)

- [Authentication and Access Control](#)
- [Dynamic Host Configuration Protocol \(DHCP\)](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [MPLS](#)
- [Multicast Protocols](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)
- [Port Security](#)
- [Routing Policy and Firewall Filters](#)
- [Software Installation and Upgrade](#)
- [User Interface and Configuration](#)
- [Virtual Chassis](#)

Authentication and Access Control

- On EX9200 Virtual Chassis, MAC address learning might fail on an authenticated interface assigned to the voice VLAN by dynamic VLAN assignment in single-secure mode. [PR1212826](#)
- On EX4200 switches, in some scenarios, the thirty-sixth port in a captive portal configuration is not redirecting to the URL as configured. This problem is seen with **set system services web-management https system-generated-certificate** configured. [PR1217743](#)
- On EX9200 switches, a MAC address corresponding to an authenticated session (dot1x) might age out as soon as traffic is not received from this MAC address for more than a few seconds (approximately 10 seconds). This leads to deletion of the authenticated session and a corresponding traffic loss. As a workaround, you can prevent the session deletion by configuring the **no-mac-binding** statement on the dot1x configuration:

```
protocols dot1x authenticator {  
    no-mac-table-binding;
```

```
}
```

PR1233261

- On an EX4300 switch or Virtual Chassis with 802.1X (dot1x) enabled, in a scenario with more than 254 clients (supplicants), many of the clients might be going to the server-reject VLAN and have limited access to the server-reject VLAN although the clients have correct credentials. For a few authenticated clients, the authentication method might be displayed as **Server-Reject** although the client was authenticated in the data VLAN. [PR1251530](#)
- On EX3300 switches, an AUTHD core file is created every time with authentication. [PR1241326](#)

Dynamic Host Configuration Protocol (DHCP)

- On EX4300 switches with DHCP relay configured, DHCP return packets—for example, DHCPREPLY and DHCPOFFER—that are received across a GRE tunnel might not be forwarded to clients, which can impact DHCP services. [PR1226868](#)
- On EX4200 Virtual Chassis, if **dhcp-relay** under **forwarding-options helpers** is configured along with **bpdu-block** and an interface configured with **bpdu-block** receives a BPDU and the interface is disabled and reenabled, a memory allocation issue may be seen that can lead to a memory exhaustion issue for DHCP relay. [PR1259918](#)

High Availability (HA) and Resiliency

- On EX4300 and QFX Series Virtual Chassis, when a switchover with GRES enabled is performed, this warning might appear: **All Packet Forwarding Engines are not ready for RE switchover and may be reset.** [PR1158881](#)
- On EX4600, QFX3500, and QFX5100 Virtual Chassis, VRRP might be preempted in case of a priority tie, but functionality is not impacted. [PR1204969](#)

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 EX4300 switches, if you configure a firewall filter on a loopback (lo0) interface to accept BGP flow and an OTHER term with the **discard** action, and the receiving host-inbound traffic with a designated TCP port 179 to the Routing Engine, existing BGP sessions might go down. [PR1090033](#)
- If you use the **request system snapshot slice alternate** command on EX2200 and EX3300 switches, a timeout error might occur and prevent completion of the file copy. The error message **error: timeout waiting for response from fpc0** is displayed when the timeout value expires before the files are copied. [PR1229520](#)
- When you load and commit a configuration on an EX2200 or EX3300 switch, the system might automatically go into db mode. As a result, you might not be able to access the switch through SSH, and a vmcore file is generated. [PR1237559](#)

- On EX4500 Virtual Chassis, there is a busy condition where the device reports incorrectly that PIC 3 has been removed. As PIC 3 is not hot-swappable, this condition should not be allowed. If this situation arises, then the device attempts to clear this illegal state by crashing chassisd. [PR1238981](#)
- EX Series switches running the ESWD process might not learn MAC addresses after a reboot if a duplicate Interface index is seen. The **show ethernet-switching interfaces detail | match Index** command can be used to confirm if each interface is showing a duplicate Interface index or if the same index is provided to two different ports. This issue is seen intermittently after a reboot when the count of **Number of VLANs * Number of Ports carrying VLANs** is in multiples of thousands. [PR1248051](#)
- On EX2200-C switches, the switch might show the Failed state for an item when you issue the **show chassis environment** operational command. This issue does not have service or traffic impact. [PR1255421](#)

Interfaces and Chassis

- On EX4300 switches, multicast traffic might be dropped after an IGMP join is received on an MC-LAG interface. [PR1167651](#)
- On EX Series Virtual Chassis that support PoE, when the master Routing Engine member is rebooted, PoE devices connected to the master might not come back online after the reboot. As a workaround, when configuring PoE interfaces, use the **set poe interface all** configuration command instead of configuring specific interfaces individually. To recover connections after seeing this issue, disable and reenab the ports with the issue. [PR1203880](#)

MPLS

- If an EX9200 switch is configured as a PE router connected to a multihomed site in an EVPN/MPLS network, RPD core files might be created on the EX9200 when more than 255 logical interfaces from the same physical interface/ESI are added to the virtual switch instance configuration. Then some logical interfaces are removed from the ESI (that is, rollback of the configuration). [PR1251473](#)

Multicast Protocols

- IGMP snooping is for IPv4 and should not affect IPv6 multicast traffic. On EX4300, EX4600, and QFX5100 switches in a Virtual Chassis configuration, IPv6 multicast packets might be affected and not be flooded in a VLAN if IGMP snooping is enabled and the ingress interface is on a different FPC than the egress interface. [PR1205416](#)
- On EX3300 and EX8200 switches, IGMP-snooping host routes might be retained after IGMP snooping has been deactivated. [PR1231751](#)

Network Management and Monitoring

- On EX4300 switches with sFlow configured, some harmless log messages regarding sFlow might be seen continuously. [PR1116568](#)
- Despite the EX4300 switch or the QFX5100 switch being configured with the network analytics feature, the analytics process might not run. As a result, the network analytics

feature might be unable to collect traffic, queue statistics, and generate reports.

[PR1165768](#), [PR1184720](#)

- On EX4600 switches, when temperatures for FPCs are polled, the temperatures might not be polled for all SNMP members. [PR1232911](#)

Platform and Infrastructure

- On an EX4300 switch, aggregated Ethernet interfaces do not display statistics for logical interfaces. [PR984998](#)
- On an EX4300 switch with Bidirectional Forwarding Detection (BFD) configured, the BFD packets might be forwarded to the best-effort queue (queue 0) instead of to the network-control queue (queue 3). When queue 0 is congested, the BFD session might flap continuously. [PR1032137](#)
- On EX4300 switches and EX4300 Virtual Chassis, PIM register messages are not forwarded to a rendezvous point (RP) when the RP is not directly connected to the first-hop router of the multicast source. [PR1134235](#)
- An EX4300 switch might drop packets received on a Layer 2 interface (for example, **set interfaces ge-1/0/24 unit 0 family ethernet-switching**) under the following conditions: (1) The interface is divided into one or more Layer 3 subinterfaces (for example, **set interfaces ge-1/0/24 unit 30 family inet address 10.0.0.254/24**). (2) The destination MAC address in the packet matches the MAC address of the Layer 3 subinterface in the routing table and in MY STATION TCAM. [PR1157058](#)
- On an EX4300 Virtual Chassis with Q-in-Q enabled, when **vlan-id-list** is configured on a C-VLAN interface and, for example, if the VLAN range vlist element is in [1-3] or [5-50], C-VLAN traffic is not sent properly across the Q-in-Q network from the C-VLAN interface. [PR1159854](#)
- On EX4300 switches with IGMP snooping enabled with **flexible-vlan-tagging** configured on ingress and egress interfaces for passthrough multicast traffic, IGMPv2 membership report messages might not be forwarded from the receiver to the sender. [PR1175954](#)
- On EX4300 switches and EX4300 Virtual Chassis, Hot Standby Router Protocol (HSRP) packets might be dropped in a VLAN if IGMP snooping is configured. As a workaround, configure the switch to flood multicast 224.0.0.2. [PR1211440](#)
- On an EX4300, if you install a firewall filter with filter-based forwarding rules to multiple bind points, it might exhaust the available TCAM. In this case, the filter is deleted from all the bind points. You can work around this issue by applying the filter to the bind points with a series of commits, applying the filter to some of the bind points with each commit. [PR1214151](#)
- On EX4300 switches, EBGP packets with ttl=1 and non-EBGP packets with ttl=1, whether destined for the device or even transit traffic, go to the same queue. In the event of a heavy inflow of non-EBGP ttl=1 packets, occasionally valid EBGP packets might be dropped, causing EBGP to flap. [PR1215863](#)
- When the **set vlans vlan-name interface all** configuration is used on EX4300, EX4600, or QFX Series switches, the Junos OS device control process (dcd) might crash as this is an unsupported configuration option on these platforms. [PR1221803](#)

- On EX Series switches except EX4300, EX4600, and EX9200 switches, Over temperature SNMP traps are sent when the CPU temperature gets higher than the bad fan temperature threshold even when there are no bad fans in the chassis. [PR1226388](#)
- On EX4300 switches, if a Layer 3 interface receives a frame with the CFI/DEI bit set to 1, this frame might be dropped and not be processed further. [PR1237945](#)
- At startup, occasionally the SFP+ ID EEPROM read fails and as a result, the SFP+ module is not recognized. As a workaround, reseal the unrecognized SFP+; for an unattended device, issue another system reboot. [PR1247172](#)
- On EX4300 switches, problems with connectivity might arise on 100M interfaces set to full duplex and half duplex or on 10M interfaces set to full duplex or half duplex. The links appear, but connectivity to end devices might not work. The port does not transmit packets even though port statistics show packets as transmitted. As a workaround: (1) Move the device to a different port. (2) Set the port to negotiate and connect a device that will autonegotiate to 1 G, full duplex; then reset the port to 10/100 full duplex or half duplex and reconnect the device. (3) Restart the pfx process. [PR1249170](#)

Port Security

- 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](#)
- High CPU caused by fxpc can lead to MACsec session drops. [PR1247479](#)

Routing Policy and Firewall Filters

- On EX Series switches other than EX9200, EX4300, and EX4600 switches, if a static MAC entry and a static ARP entry are configured, an incorrect firewall filter counter value might be displayed in command output. [PR1159940](#)
- On EX8200 Virtual Chassis, if you configure scaled firewall filters and if total terms with its own match conditions across all these filters exceed TCAM space, and you configure **examine-dhcp**, traffic will drop. [PR1215704](#)
- On EX9200 switches, if a firewall filter that has action **tcp-reset** is applied to an IRB interface, action **tcp-reset** does not work properly. [PR1219953](#)

Software Installation and Upgrade

- 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](#)

User Interface and Configuration

- On an EX Series switch that is supporting the zeroize feature, after the switch is booted up from **request system zeroize** and then a configuration is saved, the saved configuration won't be restored after the switch is rebooted. [PR1228274](#)

Virtual Chassis

- On EX4300 Virtual Chassis, a message such as **/kernel: %KERN-5: tcp_timer_keep: Dropping socket connection due to keepalive timer expiration** might be seen repeatedly. There is no service impact from the condition that causes the message (a Packet Forwarding Engine timeout trying to connect to a process that is not active). As a workaround, you can use a system-logging (syslog) filter to mask the messages. [PR1209847](#)
- On member switches in an EX Series Virtual Chassis, the **request virtual-chassis vc-port set** CLI command allows specifying an invalid or nonexistent Virtual Chassis port (VCP) interface name. An entry with the invalid VCP interface name is added to the database, and the CLI command **show virtual-chassis vc-port** displays these entries with the invalid VCP interface names, but these entries cannot subsequently be removed. [PR1215004](#)
- OID jnxFruState disappears after one of the members of the Virtual Chassis is rebooted on EX2200, EX3300, EX4200, EX4500, or EX4550 Virtual Chassis. [PR1221943](#)

Resolved Issues: Release 15.1R5

- [Authentication and Access Control](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Layer 3 Features](#)
- [MPLS](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)
- [Spanning-Tree Protocols](#)

Authentication and Access Control

- On EX4200 and EX4300 switches, dot1x server fail might not work as expected. [PR1147894](#)
- On EX9200 and EX4300 switches, 802.1X supplicants might not be reauthenticated by server fail fallback authentication after the server becomes reachable. [PR1157032](#)

- On EX9200 switches, captive portal services might not work on a switch running under Junos OS Release 15.1R4. [PR1191640](#)
- On EX4300 and EX9200 switches, dot1x scenarios involving the single-suplicant mode, mac-radius, and the server-fail deny or no server-fail action is configured, the supplicant authentication sessions might not recover after the Quiet While timer expires, once it enters the Held state. As a restoration workaround, disable and enable the interface to bring the authentication session back to the Connecting state. [PR1193944](#)

Infrastructure

- On EX8200 switches, the pfem process might crash and generate a core file. This might impact traffic. [PR1138059](#)
- On QFX5100 and EX4600 switches, in a rare timing condition, if there was already a request to gather some info from the QSFP and remove it at the same time, the packet forwarding engine manager (fxpc) might crash. [PR1151295](#)
- On EX2200-C switches, during a software upgrade to Junos OS Release 14.1X53-D35 or 15.1R3, the error messages **Triggering freezing circuitry** or **Triggering overheat circuitry** might be generated after rebooting, and then the switch shuts down. [PR1183631](#)
- On an EX8200 Virtual Chassis, doing Routing Engine failovers before booting up the line cards might cause the VLAN interface MAC address to be automatically and incorrectly set to **00:00:00:00:00:01**. [PR1185678](#)
- On EX4300, EX4600, QFX3500, QFX3600, or QFX5100 switches with **vlan-rewrite** configured on an AE interface, a VLAN rewrite might fail and result in traffic loss. [PR1186821](#)
- On EX9200 switches, periodic packet management (PPM) core files might be generated following a commit. This happens only on a large-scale setup, when the logical interface number of PFE exceeds 64. [PR1187104](#)
- On EX4200 Virtual Chassis, when an interface flaps and it has **hold-time up** configured over a long period of time (for example, 16 days), a chassis manager (chassism) process memory leak might occur due to the incorrectly accumulated task timer. About 128 bytes of the process leak every time the memory leak is triggered. [PR1188403](#)
- On EX4300 switches, VLAN rewrite does not work on aggregated links. [PR1194585](#)
- On an EX4600 switch, when you remove the 40GBASE-ER4 QSFP+ module, the **show chassis hardware** command still shows that the module is inserted. [PR1208805](#)
- On EX4200 switches and Virtual Chassis, firewall filters with syslog might not work, because as part of packet processing, packets were incorrectly mapped to the pcmd queue instead of the DFW queue. [PR1208491](#)
- On EX4200 Virtual Chassis or EX4500 or EX4550 Virtual Chassis, the Packet Forwarding Engine might not update learned MACs to an RTG active interface after RTG failover. This issue is seen with RTGs that are configured across FPCs in a Virtual Chassis setup. [PR1208491](#)
- On EX2200-C switches, the alarm Major Management Ethernet Link Down is not properly generated in cases of management link failure. [PR1209323](#)

Interfaces and Chassis

- If an EX4550-32F switch in a Virtual Chassis reboots and comes online, LACP interfaces on any member of the Virtual Chassis might go down and not come up. [PR1035280](#)
- On EX Series switches except EX9200, EX4300, and EX4600, if PoE is configured, when one IP phone is connected with a PoE interface, the phone cannot receive PoE power from the switch. [PR1174025](#)
- PoE might not work on all EX4300 ports on a mixed-mode Virtual Chassis (mixed-mode EX4600 and EX4300 or mixed-mode QFX5100 and EX4300). [PR1195946](#)
- On EX4200 and EX4550 switches on which you can configure mdi-mode manually the mode does not work properly with 15.1 releases. [PR1216549](#)

Layer 2 Features

- If an EX2200 switch is configured as a part of an ERPS ring, deactivating or deleting the ERPS configuration might cause traffic to stop forwarding through one or more VLANs. [PR1189585](#)
- An EX Series switch might not process ERPS PDUs that are received from other nodes. This could lead to the ERPS ring not operating correctly. [PR1190007](#)
- On EX4300 Virtual Chassis, a Layer 2 interface might not be associated with the default VLAN after you add the interface to the ethernet-switching family. [PR1192679](#)
- On EX9200, EX4300, EX4600, QFX3500, QFX3600, QFX3500, and QFX5100 switches, if 'set protocols xstp interface all edge' is configured in combination with 'set protocols xstp bpdu-block-on-edge', interfaces do not go down (Disabled - Bpdu-Inconsistent) when they receive BPDUs; they transition to non-edge. If an interface is configured specifically with 'set protocols xstp interface interface-name edge', then when that interface receives a BPDU, it goes down or transitions into Disabled - Bpdu-Inconsistent correctly. As a workaround, configure **set protocols layer2-control bpdu-block interface all**. [PR1210678](#)

Layer 3 Features

- On a switch that has **secure-access-port** configured, when you change the MTU size of interfaces and commit, VRRP sessions might flap between the VRRP master and backup. [PR1163652](#)
- On EX2200, EX3300, EX4200, EX4500, EX4550, EX6200, and EX8200 switches, when VRRP configuration changes from ethernet-switching to inet family and vice-versa, then the local IP of the master VRRP switch cannot be reached on the backup VRRP switch and vice-versa. Virtual IP is always reached on both switches. [PR1171220](#)

MPLS

- On EX4600 switches, when traffic enters an MPLS interface and is destined to the loopback interface in the routing instance, the firewall filter might not work properly. [PR1205626](#)

Platform and Infrastructure

- If you use the **load replace** command or the **load merge** command to configure a device and have included an annotation just before a **delete** action in the loaded configuration file, the management daemon (mgd) might create a core file. [PR1064036](#)
- On EX4300 Virtual Chassis, if a Q-in-Q S-VLAN interface with MC-LAG is configured, when a backup EX4300 is acting as master, the connection to the management IP address through the interface might be lost, causing a management traffic loss. [PR1131755](#)
- On EX4300 switches, when xSTP is configured, if you unplug and then plug in one loopback cable between ports of different FPCs, an interface might go down and a BPDU error might be detected on this port, causing traffic to drop on another egress port. [PR1160114](#)
- On EX4300 switches, when DHCP security is enabled on a VLAN, unicast packets (for example, DHCP Offers and ACKs) might be forwarded to all ports in the VLAN. [PR1172730](#)
- On EX4300 switches, if an Ethernet port receives a frame with a CFI/DEI bit set to 1, then this frame would not be bridged to an untagged (access) port; it could be bridged to a trunk port. [PR1176770](#)
- When IGMP snooping and storm control are enabled, EX Series switches are supposed to forward traffic with destination IP address 224.0.0.0/24 to all ports on a VLAN. But for EX4300, except for the well-known addresses in this range—for example, 224.0.0.5/6 for OSPF, 224.0.0.20 for VRRP—all other multicast traffic with a destination in 224.0.0.0/24 is dropped. [PR1176802](#)
- If you upgrade the Power over Ethernet (PoE) firmware on a member of an EX4300 Virtual Chassis, the PoE firmware upgrade process might fail or get interrupted on that member switch. You can recognize that this problem has occurred if the member switch is not listed in the command output when you issue the "show poe controller" command. The problem is also indicated if you issue the ?show chassis firmware detail? command and the ?PoE firmware? version field is not shown in the output or has a value of 0.0.0.0. [PR1178780](#)
- On EX4300 switches, if there is a mismatch in the speed configuration between two interfaces, the link might be autonegotiated to half-duplex mode instead of full-duplex mode. [PR1183043](#)
- On EX4300 switches configured with dscp and 802.1p rewrite rules on an interface, if you delete 802.1p rewrite-rules from the interface, the 802.1p rewrite might still happen along with the dscp rewrite. [PR1187175](#)
- On EX4300, EX4600, and QFX Series switches with VSTP enabled for multiple VLANs and participated in a VSTP topology, when BPDU packets are received on the Packet

Forwarding Engine from other switches, the switch sends BPDU packets to the Routing Engine for further VSTP computing. But, in rare cases, the switch might not send VSTP packets for all VLANs to the Routing Engine. For example, for a VLAN, BPDU packets are not reaching the Routing Engine, even though VSTP is enabled for that VLAN. This will result in this VLAN considering itself the root bridge and advertising itself as the root bridge and sending BPDUs to other VSTP switches. Other switches might block related ports. [PR1187499](#)

- On EX Series Virtual Chassis, a next-hop change message might not be sent from the Routing Engine when a LAG member is added or deleted, and hence packets are dropped in the Packet Forwarding Engine, as the next hop is not updated properly. [PR1201740](#)
- When seating an SFP in a operating EX4300 switch, sometimes the SFP would be recognized as unsupported or as an SFP+-10G. The cause is that the switch reads the EEPROM information of the SFP before waiting long enough for SFP initialization. [PR1202730](#)
- On EX4300 switches, if you activate DHCP security features for IPv6, a JDHCPD core file might be generated. [PR1212239](#)
- On an EX9200 switch, with a services REST configuration, after a reboot, the configuration is not applied and SSH stops working. [PR1212425](#)
- 1G fiber link ports might be down with MACsec configured on EX4300 switches when the switch is rebooted. [PR1172833](#)

Routing Protocols

- On EX4300 Virtual Chassis with IGMP snooping enabled, when IGMP hosts subscribe to the same group, IGMP queries might not go through between a member in the linecard role and the master. [PR1200008](#)

Spanning-Tree Protocols

- On EX Series switches except for EX4300, EX4600, and EX9200, while the switch is processing an xSTP-disabled interface with a BPDU block configuration, current code flow might set the bpdv_control flag for RSTP-enabled interfaces as well. This might result in RSTP-enabled ports becoming blocked when they receive a BPDU. [PR1185402](#)
- On EX9200, EX4300, EX4600, QFX3500, QFX3600, and QFX5100 platforms, when any type of spanning tree (STP, RSTP, MSTP, or VSTP) is configured, the MAC address part of the bridge ID might be set to all zeros (for example, 4096.00:00:00:00:00:00) after you power cycle the device without issuing the **request system halt** command. As a workaround, issue the **restart l2-learning** command. [PR1201493](#)

Resolved Issues: Release 15.1R4

- [High Availability \(HA\) and Resiliency](#)
- [Interfaces and Chassis](#)
- [Network Management and Monitoring](#)
- [Platform and Infrastructure](#)

- [Software Installation and Upgrade](#)
- [Spanning Tree Protocols](#)
- [User Interface and Configuration](#)
- [Virtual Chassis and Virtual Chassis Fabric \(VCF\)](#)

High Availability (HA) and Resiliency

- On EX4300 Virtual Chassis, after a nonstop software upgrade (NSSU), the master might detect the backup coming up after the upgrade and reprogram the trunk, even though the backup member links are down. Traffic might drop when the master tries to push the traffic through trunk members that have not yet come up. Traffic resumes after the links come up. [PR1115398](#)
- On EX4300 Virtual Chassis, traffic loss might occur for about 10 seconds when the master leaves the Virtual Chassis for upgrade. [PR1173754](#)

Interfaces and Chassis

- On EX2200 switches, in Ethernet ring protection switching (ERPS) configurations, no VLAN is included in **data-channel** if **data-channel** is not explicitly configured, and a MAC flush does not happen for any data VLAN while the switch receives an SF signal, which might cause a traffic issue before the MAC address ages out. [PR1152188](#)
- On EX2200 switches, in an ERPS configuration, many SF (signal failure) packets might appear in a link-end ring node during a link failure that existed for a short time. [PR1169372](#)
- On EX4300 Virtual Chassis, Layer 2 multicast might not work properly when both Layer 2 and Layer 3 entries are present for the same group on two different integrated routing and bridging (IRB) interfaces. [PR1183531](#)

Network Management and Monitoring

- On EX9200 switches, ingress sFlow samples of packets routed on an integrated routing and bridging (IRB) interface might be dropped. [PR1147719](#)
- On EX9200 switches, an sFlow flow sample with an incorrect frame length value in a raw packet header might be generated for frames larger than 128 bytes, and traffic volumes calculated based on frame length and sampling rate values in the sFlow collector might be inaccurate. [PR1152275](#)
- On EX9200 switches, eventd might run out of memory and crash because of excessive kernel logging. [PR1162722](#)

Platform and Infrastructure

- On EX4500, EX4550, EX6200, and EX8200 switches, if you replace a 1-gigabit SFP transceiver with a 10-gigabit SFP+ transceiver on one port, the adjacent port might go down. For example, if you install an SFP transceiver in each of port-0/0/36 and port-0/0/37, and replace each SFP transceiver with an SFP+ transceiver in port-0/0/36 and port-0/0/37, then port-0/0/36 might go down during the insertion of the SFP+ transceiver in port-0/0/37. [PR1073184](#)
- In an EX8200 Virtual Chassis in which the external Routing Engine (XRE200) has two DC power supplies installed, when one power supply fails, no logs or SNMP traps are generated. [PR1162165](#)
- If a configuration is pushed to an EX Series switch using Zero Touch Provisioning (ZTP), then after a subsequent reboot, the configuration might be deleted. [PR1170165](#)
- On EX3300 and EX4200 switches, after the **request system zeroize media** command has been executed, J-Web might stop responding. [PR1177214](#)
- On an EX4300 switch or Virtual Chassis, the chassisd daemon might get stuck and become unresponsive. If you issue a chassisd-related show command, the command returns the error message **error: the chassis-control subsystem is not responding to management requests**. [PR1038830](#)
- On ARM platforms such as EX3300 switches, configuring internal IPsec security associations containing the authentication hmac-sha2-256 might throw a kernel alignment exception. [PR1149565](#)
- On EX4300 switches, if IGMP snooping is enabled, packets with destination 224.0.0.0/24 might be dropped, except for well-known addresses (for example, 224.0.0.5/6 for OSPF). [PR1167859](#)
- On EX4300 switches, ICMP-tagged packets might transit the egress interface of a PVLAN access port. [PR1169116](#)

Software Installation and Upgrade

- On EX8200 Virtual Chassis, traffic might be lost for multicast and Layer 3 protocols (such as RIP, OSPF, BGP, and VRRP) during a nonstop software upgrade (NSSU). [PR1185456](#)
- On EX6200 switches, multicast traffic and Layer 3 protocol traffic (such as RIP, OSPF, BGP, and VRRP) might be lost during a nonstop software upgrade (NSSU). [PR1185816](#)
- On EX8200 switches, multicast traffic might be lost during a nonstop software upgrade (NSSU). [PR1185888](#)

Spanning Tree Protocols

- On EX4300, EX4600, and EX9200 switches, when root guard is in effect or cleared, there appropriate system log messages might not be displayed. [PR1176240](#)

User Interface and Configuration

- On a device configured with an SSH public key for which the string buffer size exceeds 1 Kb, if you load the configuration by using the **load override** command, the management daemon (mgd) might create a core file. [PR1153392](#)

Virtual Chassis and Virtual Chassis Fabric (VCF)

- On EX3300 Virtual Chassis, the **vcp-snmp-statistics** configuration statement is not listed in the **[edit virtual-chassis]** hierarchy. [PR1178467](#)

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 375: 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

- Setting link speed to 100 Mbps does not work in the following situations:
 - When network interfaces are used on an EX4600 switch
 - When an EX4600-EM-8F expansion module is installed in a QFX5100-24Q switch or an EX4600 switch

[PR1032257](#)

- 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 bpdv-block 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. [PR111669](#)

Resolved Issues: Release 15.1R2

- [Class of Service \(CoS\)](#)
- [Dynamic Host Configuration Protocol](#)
- [Infrastructure](#)
- [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 EX9200 switches, a DHCPv6 security dynamic entry binding might not work as expected, resulting in the DHCPv6 bindings being stuck in the wait state. [PR1092885](#)
- 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](#)

Infrastructure

- Unnecessary `fpc0 dfw_counter_get_by_name failed inst 0 policer index 0 status 7` log messages are seen when either `show firewall counter` or `snmp mib get jnxFirewallCounterTable` is executed. [PR1035113](#)

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 **family ethernet-switching** is configured on an interface that is also configured with **encapsulation extended-vlan-bridge**, then transit packets (for example, IP, ping, or Q-in-Q packets) might be dropped on this interface. [PR1078076](#)
- 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](#)

Resolved Issues: Release 15.1R1

- [Interfaces and Chassis](#)

Interfaces and Chassis

- On EX Series switches on which Link Aggregation Control Protocol (LACP) is enabled on a link aggregation group (LAG) interface, after you reboot the master Routing Engine and if the first LACP packet is dropped during switchover, LACP might get stuck in the same state for a long time (about 10 seconds), causing the LAG interface to flap and traffic drop on the LAG interface. [PR976213](#)

- See Also**
- [New and Changed Features on page 37](#)
 - [Changes in Behavior and Syntax on page 45](#)
 - [Known Behavior on page 46](#)
 - [Known Issues on page 53](#)
 - [Documentation Updates on page 76](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 77](#)
 - [Product Compatibility on page 78](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R6 for the EX Series switches documentation.

- [Changes to the Junos OS for EX Series Documentation on page 77](#)
- [Errata in the Junos OS for EX Series Documentation on page 77](#)

Changes to the Junos OS for EX Series Documentation

Network Interfaces Feature Guide for EX4300 Switches

- Half-duplex link support has been added to the EX4300 switch starting with Junos OS Release 15.1R4. The *Network Interfaces Feature Guide for EX4300 Switches* has not yet been updated to show this support. See the description of this feature in [“New and Changed Features” on page 37](#).

Errata in the Junos OS for EX Series Documentation

Junos OS Release 15.1 Release Notes

- The EX3200 switch is not supported in Junos OS Release 15.1. We have removed references to EX3200 switches in *Junos OS Release 15.1* release notes, but note that PDF versions of the release notes that you have downloaded or saved might not reflect those updates.
- PR976213 was resolved in Junos OS Release 15.1R1 but was erroneously listed in the Known Behavior of the *Junos OS Release 15.1* release notes. We have moved the PR to Resolved Issues in the release notes, but note that PDF versions of the release notes might not reflect that update.

- See Also**
- [New and Changed Features on page 37](#)
 - [Changes in Behavior and Syntax on page 45](#)
 - [Known Behavior on page 46](#)
 - [Known Issues on page 53](#)
 - [Resolved Issues on page 54](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 77](#)
 - [Product Compatibility on page 78](#)

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 77](#)

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 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

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

- See Also**
- [New and Changed Features on page 37](#)
 - [Changes in Behavior and Syntax on page 45](#)
 - [Known Behavior on page 46](#)
 - [Known Issues on page 53](#)
 - [Resolved Issues on page 54](#)
 - [Documentation Updates on page 76](#)
 - [Product Compatibility on page 78](#)

Product Compatibility

- [Hardware Compatibility on page 78](#)

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 <https://pathfinder.juniper.net/feature-explorer/>.

- See Also**
- [New and Changed Features on page 37](#)
 - [Changes in Behavior and Syntax on page 45](#)
 - [Known Behavior on page 46](#)

- [Known Issues on page 53](#)
- [Resolved Issues on page 54](#)
- [Documentation Updates on page 76](#)
- [Migration, Upgrade, and Downgrade Instructions on page 77](#)

Junos OS Release Notes for M Series Multiservice Edge Routers, MX Series 5G Universal Routing Platforms, and T Series Core Routers

These release notes accompany Junos OS Release 15.1R6 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 https://www.juniper.net/documentation/product/en_US/junos-os.

- [New and Changed Features on page 80](#)
- [Changes in Behavior and Syntax on page 131](#)
- [Known Behavior on page 165](#)
- [Known Issues on page 170](#)
- [Resolved Issues on page 192](#)
- [Documentation Updates on page 312](#)
- [Migration, Upgrade, and Downgrade Instructions on page 320](#)
- [Product Compatibility on page 329](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R6 for the M Series, MX Series, and T Series.

- [Hardware on page 81](#)
- [Bridging and Learning on page 81](#)
- [Class of Service \(CoS\) on page 82](#)
- [High Availability \(HA\) and Resiliency on page 83](#)
- [Interfaces and Chassis on page 86](#)
- [IPv6 on page 90](#)
- [Junos OS XML API and Scripting on page 90](#)
- [Layer 2 Features on page 91](#)
- [Management on page 93](#)
- [MPLS on page 94](#)
- [Multicast on page 96](#)
- [Network Management and Monitoring on page 97](#)
- [Routing Policy and Firewall Filters on page 99](#)
- [Routing Protocols on page 100](#)
- [Services Applications on page 102](#)
- [Software Defined Networking on page 108](#)
- [Software Installation and Upgrade on page 109](#)
- [Software Licensing on page 109](#)

- [Subscriber Management and Services \(MX Series\) on page 112](#)
- [System Logging on page 129](#)
- [User Interface and Configuration on page 129](#)
- [VPNs on page 129](#)

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](#).]

- **Support for suppressing the default classifier (MX Series)**—Beginning with Junos OS Release 16.1R1, you can disable the application of the default classifier on an interface or a routing instance to preserve the incoming classifier. This is done by applying the **no-default** option at the **[edit class-of-service routing-instances routing-instance-name classifiers]** hierarchy level. This is useful, for example, in a bridge domain, where the default classifier for the interface overrides the configured classifier for the domain.

[See [Applying Behavior Aggregate Classifiers to Logical Interfaces](#).]

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](#).]

- **New option providing detailed information about the stages in an ISSU operation (MX240, MX480, MX960, MX2010, and MX2020 Universal Routing Platforms)**—Starting in Junos OS Release 15.1R6 a new option, **verbose**, is added at the **[request system software in-service-upgrade package-name]** hierarchy level. This new option provides information about all stages in a unified ISSU operation.

When this option added to the command, the output provides the following additional information:

- State of various Routing Engine daemons
- Daemon that caused the failure of a unified ISSU operation in cases where the ISSU is aborted because of daemon readiness check failure
- Progress of a unified ISSU operation, such as the initialization, hardware synchronization, and completion

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](#).]

- **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 2 km reach over single-mode fiber, 100 m (with OM3 MMF cable), or 150 m (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, 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](#).]

- **Enhancement to policer configuration**—Starting in Junos OS Release 15.1R6, you can configure the MPC to take a value in the range 0 through 5 for the policer tick byte by using the **policer-limit** statement at the **[edit chassis]** hierarchy level. If this statement is not configured, the policer tick byte can take values up to 7, which is the default behavior. You can use the **set chassis policer-limit** command to enable this feature.

You must restart the MPC or the router for the changes to take effect.

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 [Junos OS Layer 2 Switching and Bridging Library](#).]

- **Hot-standby support for VPLS redundant pseudowires**—Starting in Release 15.1R4, Junos OS enables you to configure redundant pseudowires. If a primary pseudowire fails, Junos OS switches service to a preconfigured redundant pseudowire.

The time required for the redundant pseudowire to recover traffic from the primary pseudowire depends on the number of pseudowires and the option configured for pseudowire redundancy. There are three options:

- Backup redundancy
- Standby redundancy
- Hot-standby

The hot-standby option enables Junos OS to reduce the amount of traffic it discards during a transition from a primary to redundant pseudowire. Both the active and standby paths are kept open within the Layer 2 domain. Now you can configure the hot-standby option to configure pseudowires for virtual private LAN services (VPLS) running the Label Distribution Protocol (LDP).

- **Implicit maximum bandwidth for inline services for L2TP LNS (MX Series)**—Starting in Junos OS Release 15.1R5, you are no longer required to explicitly specify a bandwidth for L2TP LNS tunnel traffic using inline services. When you do not specify a bandwidth, the maximum bandwidth supported on the PIC is automatically available for the inline services; inline services can use up to this maximum value. For example:

```
user@host# set chassis fpc 3 pic 0 inline-services
user@host# set chassis fpc 3 pic 1 inline-services
```

```
user@host> show interfaces si-3/0/0
```

```
Physical interface: si-3/0/0, Enabled, Physical link is Up
Interface index: 181, SNMP ifIndex: 561
Type: Adaptive-Services, Link-level type: Adaptive-Services,
MTU: 9192, Speed: 100000mbps
...
```

```
user@host> show interfaces si-3/1/0
```

```
Physical interface: si-3/1/0, Enabled, Physical link is Up
Interface index: 182, SNMP ifIndex: 562
Type: Adaptive-Services, Link-level type: Adaptive-Services,
MTU: 9192, Speed: 100000mbps
...
```

In earlier releases, you must specify a bandwidth to enable inline services by including the **bandwidth** statement with the **inline-services** statement.

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 [Dynamic Bandwidth Management Using Container LSP Overview](#).]

- **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 `rpdpd` 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 DMEM 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 15.1R1, 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.

[See [RSVP LSP Tunnels Overview](#)]

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.](#)]

Starting with Junos OS Release 15.1R1, you can configure the interfaces for routing instances at the `[edit routing-instances instance-name protocols mpls]`, `[edit routing-instances instance-name protocols rsdp]`, and `[edit routing-instances instance-name protocols ldp]` hierarchy levels. You cannot configure an interface that already exists in a routing instance under `[edit protocols mpls]`, `[edit protocols rsdp]`, and `[edit protocols ldp]` hierarchy levels.

- **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 IPv6 neighbor discovery protocol (NDP) cache protection on a per-interface basis. NDP performs address resolution and maintains the neighbor cache, and can be susceptible to denial-of-service (DoS) attacks that overwhelm the device's control plane with unassigned address resolution requests, resulting in a cache overflow. One strategy for mitigating this type of DoS attack is to enforce neighbor discovery queue limits, restricting the overall number of IPV6 neighbors and new unresolved next-hop addresses that can be added to the cache.

The device has default cache limits that can be changed system-wide, or you can use this feature to override default or system-wide limits on a per-interface basis.

[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 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 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 **elf-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 H.323 NAT on MS-MPC and MS-MIC (MX Series routers)**—Starting in Junos OS Release 15.1R5, the H.323 ALG is supported in NAT-44 rules and IPv4 stateful-firewall rules on the MX Series. H.323 is a legacy VoIP protocol.

To configure H.323 in a NAT-44 rule, include the **application-sets junos-h323-suite** statement at the **[edit services nat rule *rule-name* term *term-name* from]** hierarchy level. To configure H.323 in a stateful-firewall rule, include the **application-sets junos-h323-suite** statement at the **[edit services stateful-firewall rule *rule-name* term *term-name* from]** hierarchy level.

To show H.323 ALG statistics, issue the **show services alg statistics application-protocol h323** command.

- **Class-of-service (CoS) marking and reclassification for the MS-MICs and MS-MPCs**—Starting with Junos Release 15.1R5, the MS-MIC and MS-MPC support CoS configuration, which enables you to configure Differentiated Services code point (DSCP) marking and forwarding-class assignment for packets transiting the MS-MIC or MS-MPC. You can configure the CoS service alongside the stateful firewall and NAT services, using a similar rule structure.

[See [Configuring CoS Rules](#).]

- **Support for IKE and IPsec on NAPT-44 and NAT64 (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1R5, you can enable the passing of IKE and IPsec packets through NAPT-44 and NAT64 filters between IPsec peers that are not NAT-T compliant by using the IKE-ESP-TUNNEL-MODE-NAT-ALG on MS-MPCs and MS-MICs.

Use the following hierarchy to enable the IKE-ESP-TUNNEL-MODE-NAT-ALG:

```
[edit applications]
application ike-esp-application-name {
  application-protocol ike-esp-nat;
  protocol udp;
  destination-port 500;
  inactivity-timeout 3600;
}
application-set ike-esp-application-set-name {
  application ike-esp-application-name;
}

[edit services nat]
pool ike-isp-nat-pool-name {
  address ip-prefix;
  port automatic;
}
rule rule-name {
  match-direction input;
  term 0 {
    from {
      source-address address;
      application-sets ike-esp-application-set-name;
    }
    then {
      translated {
        source-pool ike-isp-nat-pool-name;
        translation-type napt-44;
      }
    }
  }
}
```

- **Support for AMS warm standby on MS-MPC and MS-MIC (MX Series routers)**—Starting in Junos OS Release 15.1R5, one service interface can be the backup interface for multiple service interfaces. This feature is called AMS warm standby. To make a service interface the backup for multiple service interfaces, you configure an AMS interface for each service interface you want to protect. Each of these AMS

interfaces has two member interfaces—a primary member interface, which is the service interface you want to protect, and the secondary member interface, which is the backup service interface. You can use the same secondary member interface in multiple AMS interfaces.

To configure a warm-standby AMS interface, include the **primary mams-a/b/0** statement and the **secondary mams-a/b/0** statement at the **[edit interfaces amsn redundancy-options]** hierarchy level.

If you use **redundancy-options** in an AMS interface, you cannot use **load-balancing-options** in the same AMS interface.

You cannot use the same member interface in both an AMS interface that includes **load-balancing-options** and an AMS interface that includes **redundancy-options**.

To show the state of an AMS interface configured with warm standby, issue the **show interfaces redundancy** command.

To switch from the primary interface to the secondary interface, issue the **request interface switchover amsn** command.

To revert to the primary interface from the secondary interface, issue the **request interface revert amsn** command.



NOTE: Support for IPv6 on RPM probes is not supported in Junos OS Release 15.1R1. However, documentation for this feature is included in the Junos OS 15.1R1 documentation set.

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.](#)]

Software Licensing

- **Licensing enhancements (M Series, MX Series and T Series)**—Starting with Junos OS Release 15.1R1, licensing enhancements on routers running Junos OS enable you to configure and delete license keys in a Junos OS CLI configuration file. The license keys are validated and installed after a successful commit of the configuration file. If a license key is invalid, the commit fails and issues an error message. You can configure individual license keys or multiple license keys by issuing Junos OS CLI commands or by loading the license key configuration contained in a file. All installed license keys are stored in the **/config/license/** directory.

To install an individual license key in the Junos OS CLI, issue the **set system license keys *key name*** command, and then issue the **commit** command.

For example:

```
[edit]
root@switch# set system license keys key "JUNOS_TEST_LIC_FEAT aeaqeb qbmqds
qwwsxe okyvou 6v57u5 zt6ie6 uwl3zh assvnu e2ptl5 soxawy vtfh7k axwnno m5w54j
6z"
root@switch# commit
commit complete
```

To verify that the license key was installed, issue the **show system license** command.

For example:

```

root@switch> show system license
License usage:

              Licenses      Licenses      Licenses      Expiry
Feature name      used      installed      needed
sdk-test-feat1      0          1          0
permanent

Licenses installed:
License identifier: JUNOS_TEST_LIC_FEAT
License version: 2
Features:
  sdk-test-feat1    - JUNOS SDK Test Feature 1
  permanent

```

To install multiple license keys in the Junos OS CLI, issue the **set system license keys** *key name* command, and then issue the **commit** command.

For example:

```

[edit]
root@switch# set system license keys key "key_1"
set system license keys key "key_2"
set system license keys key "key_2"
set system license keys key "key_4"
root@switch# commit
commit complete

```

To verify that the license key was installed, issue the **show system license** command.

To install an individual license key configuration in a file, issue the **cat** command:

For example:

```

[edit]
root@switch%cat license.conf
system {
  license {
    keys {
      key "JUNOS_TEST_LIC_FEAT aeaqeb qbmqds qwwsxe okyvou 6v57u5 zt6ie6
uw13zh assvnu e2pt15 soxawy vtfh7k axwnno m5w54j 6z";
    }
  }
}

```

Load and merge the license configuration file.

For example:

```

[edit]
root@switch# load merge license.conf
load complete

```

Issue the **show | compare** command to see the configuration, and then issue the **commit** command.

For example:

```
[edit]
root@switch# show | compare
[edit system]
+   license {
+       keys {
+           key "JUNOS_TEST_LIC_FEAT aeaqeb qbmqds qwwsxe okyvou 6v57u5
zt6ie6 uw13zh assvnu e2pt15 soxawy vtfh7k axwnno m5w54j 6z";
+       }
+   }
[edit]
root@switch# commit
```

To verify that the license key was installed, issue the **show system license** command.

For example:

```
root@switch> show system license
License usage:

Feature name          Licenses      Licenses      Licenses      Expiry
                    used       installed     needed
sdk-test-feat1        0           1             0
permanent

Licenses installed:
License identifier: JUNOS_TEST_LIC_FEAT
License version: 2
Features:
  sdk-test-feat1  - JUNOS SDK Test Feature 1
  permanent
```

To install multiple license keys in a file, issue the **cat** command:

For example:

```
[edit]
root@switch%cat license.conf
system
{
  license
  {
    keys
    {
      key "key_1"
      key "key_2"
      key "key_3"
      ...
      key "key_n"
    }
  }
}
```

Load and merge the license configuration file, and then issue the **commit** command.

For example:

```
[edit]
root@switch# load merge license.conf
load complete
```

```
[edit]
root@switch# commit
```

To verify that the license key was installed, issue the **show system license** command.

You can also delete or deactivate individual and multiple license keys in the Junos OS CLI by issuing the **delete system license keys** or **deactivate system license keys** commands. Do not use the **request system license delete** command to delete the license keys.

For example, to issue the **delete system license keys** command:

```
[edit]
root@switch# delete system license keys
root@switch# commit
```

Subscriber Management and Services (MX Series)

- **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 Junos OS 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 (300 MB 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 for subscriber management when you upgrade from a release that does not support enhanced subscriber management (Junos OS Release 14.2 or earlier) to a release that does support enhanced subscriber management (15.1R3 and later), all subscriber sessions and subscriber state are lost after such an 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.
- **Extended hardware support for L2TP inline IP reassembly (MX Series)**—Starting in Junos OS Release 15.1R3, L2TP inline IP reassembly support is extended to the following MPCs:

MPC2E-3D-NG	MPC5E-40G10G
MPC2E-3D-NG-Q	MPC5EQ-40G10G
MPC3E-3D-NG	MPC5E-100G10G
MPC3E-3D-NG-Q	MPC5EQ-100G10G

- **Monitoring only ingress traffic for subscriber idle timeouts**—Starting in Junos OS Release 15.1R3, you can specify that only ingress traffic is monitored for subscriber idle timeout processing. When you include the **client-idle-timeout-ingress-only** statement at the **[edit access-profile *profile-name* session-options]** hierarchy level, subscribers are logged out or disconnected if no ingress traffic is received for the duration of the idle timeout period. Egress traffic is not monitored. When you do not include this statement, both ingress and egress traffic are monitored during the timeout period to determine whether subscribers are logged out or disconnected.
- **Support for service session termination causes (MX Series)**—Starting in Junos OS Release 15.1R3, new internal identifiers are available that identify the reasons that authd initiates termination of individual service sessions. In earlier releases, the termination cause for a service session is the same as that for the parent subscriber session.

The service termination causes map to default code values that are reported in the RADIUS Acct-Terminate-Cause attribute (49) in Acct-Stop messages for the service. You can use the new **service-shutdown** option with the **terminate-code aaa** statement at the **[edit access]** hierarchy level to remap any of the new termination causes to any number in the range 1 through 4,294,967,295:

- **network-logout**—Termination was initiated by deactivation of one family for a dual-stack subscriber, typically triggered by termination of the corresponding Layer 3 access protocol. Default code value is 6.
- **remote-reset**—Termination was initiated by an external authority, such as a RADIUS CoA service-deactivation. Default code value is 10.

- **subscriber-logout**—Overrides the default inheritance of the subscriber session value with a different value when you map it to a different value. Default code value is 1, meaning that it inherits the terminate cause from the parent subscriber session.
- **time-limit**—Service time limit was reached. Default code value is 5.
- **volume-limit**—Service traffic volume limit was reached. Default code value is 10.

The **show network-access aaa terminate-code aaa detail** command displays the new termination causes and their current code values.

- **New option for service type added to test aaa commands (MX Series)**—Starting in Junos OS Release 15.1R4, you can include the **service-type** option with the **test aaa ppp user** and **test aaa dhcp user** commands to test the AAA configuration of a subscriber. You can use this option to distinguish a test session from an actual subscriber session. The option specifies a value for the Service-Type RADIUS attribute [6] in the test Access-Request message; when you do not include this option, the test uses a service type of Framed. You can specify a number in the range 1 through 255, or you can specify one of the following strings that corresponds to an RFC-defined service type; the numbers are the values that are carried in the RADIUS attribute to specify the service:

administrative (6)	callback-nas-prompt (9)
authenticate-only (8)	framed (2)
call-check (10)	login (1)
callback-admin (11)	nas-prompt (7)
callback-framed (4)	outbound (5)
callback-login (3)	—

When the Service-Type RADIUS attribute [6] is received in an Access-Accept message, it overrides the value inserted in the Access-Request message by this command.

- **New predefined variable for dynamic underlying interfaces (MX Series)**—Starting in Junos OS Release 15.1R4, you can use the Juniper Networks predefined variable, **\$junos-underlying-ifd-name**, to reference the underlying physical interface when you configure CoS properties for an underlying logical interface in a dynamic profile. The new variable is useful when the **\$junos-interface-ifd-name** variable already references a different physical interface, such as in configurations with stacked logical interfaces. For example, in a PPPoE session where the PPP logical interface is stacked over a demux VLAN logical interface, **\$junos-interface-ifd-name** is set to the pp0 physical interface. In this case you can specify the **\$junos-underlying-ifd-name** predefined variable with the **interfaces** statement at the **[edit dynamic-profiles profile-name class-of-service]** hierarchy level to reference the underlying physical interface.
- **Increase in range for RADIUS server accounting-retry statement (MX Series)**—Starting in Junos OS Release 15.1R4, you can configure the router to make a maximum of 60 attempts to send interim accounting messages to the RADIUS accounting server when

it has received no response. In earlier releases, the maximum number of attempts is 30.



BEST PRACTICE: We recommend that you do not configure a retry duration greater than or equal to 30 accounting retries times 90 seconds per accounting timeout period. Configure fewer retries, a shorter timeout, or both.

- **New predefined variables and Juniper Networks VSAs for family any interface filters (MX Series)**—Starting in Junos OS Release 15.1R4, you can use the `$junos-input-interface-filter` and `$junos-output-interface-filter` predefined variables to attach a filter to a dynamic interface created for family any. The filter names are derived from the Juniper Networks VSAs, Input-Interface-Filter (26-191) and Output-Interface-filter (26-192). These VSAs are conveyed in the following RADIUS messages: Access-Request, Acct-Start, Acct-Stop, and Acct-Interim-Interval. You can specify the variables as the filter names with `input` and `output` statements at the `[edit dynamic-profiles profile-name interfaces interface-name unit logical-interface-number filter]` hierarchy level.
- **New predefined variable to group subscribers on a physical interface (MX Series)**—Starting in Junos OS Release 15.1R4, you can specify the new Juniper Networks predefined variable, `$junos-phy-ifd-interface-set-name`, with the `interface-set` statement at the `[edit dynamic-profiles profile-name interfaces]` hierarchy level to configure an interface set associated with the underlying physical interface in a dynamic profile. This predefined variable enables you to group all the subscribers on a specific physical interface so that you can apply services to the entire group of subscribers.

Another use case is optimizing CoS level 2 node resources by grouping residential subscribers into an interface set associated with the physical interface in a topology where residential and business subscribers share the interface, enabling the use of CoS level 2 nodes for the interface set rather than for each residential interface.

- **Configuring default values for routing instances (MX Series)**—Starting in Junos OS Release 15.1R4, you can define a default value for the Juniper Networks predefined variable, `$junos-routing-instance`. This value is used in the event RADIUS does not supply a value for `$junos-routing-instance`. To configure a default value, use the `predefined-variable-defaults` statement at the `[edit dynamic-profiles]` hierarchy level. For example, to set the default value to RI-default:

```
[edit dynamic-profiles profile-name]
user@host# set predefined-variable-defaults routing-instance RI-default
```

- **Hot-standby support for VPLS redundant PWs**—Starting in Junos 15.1R4, Junos OS enables you to configure redundant pseudowires (PWs). If a primary PW fails, Junos OS switches service to a preconfigured redundant PW.

The time required for the redundant PW to recover traffic from the primary PW depends on the number of PWs and the option configured for PW redundancy. There are three options:

- Backup redundancy

- Standby redundancy
- Hot-standby

The hot-standby option enables Junos OS to reduce the amount of traffic it discards during a transition from a primary to redundant PW. Both the active and standby paths are kept open within the Layer 2 domain. Now you can configure the hot-standby option to configure PWs for virtual private LAN services (VPLS) running the Label Distribution Protocol (LDP).

- **Enhanced DHCP dual-stack support (MX Series)**—Starting in Junos OS Release 15.1R4, subscriber management supports a single-session DHCP dual-stack model that provides a more efficient configuration and management of dual-stack subscribers.

The single-session dual-stack model addresses session-related inefficiencies that exist in the traditional dual-stack—for example, the new model requires single sessions for authentication and accounting, as opposed to multiple sessions that are often needed in a traditional dual-stack configuration. The single-session dual-stack model also simplifies router configuration, reduces RADIUS message load, and improves accounting session performance for subscriber households with dual-stack environments.

[See [Single-Session DHCP Dual-Stack Overview](#).]

- **DHCPv6 subscriber identification criteria and automatic logout**—Starting in Junos OS Release 15.1R5, the DHCPv6 local server and the DHCPv6 relay agent can identify a DHCPv6 client by the incoming-interface option in addition to the client identifier. The incoming interface allows only one client device to connect on the interface. If the client device changes—that is, if DHCPv6 receives a Solicit message from a client whose incoming interface matches the existing interface—DHCPv6 automatically logs out the existing client without waiting for the normal lease expiration. It deletes the existing client binding and creates a binding for the newly connected device.

For DHCPv6 local server, include the **client-negotiation-match incoming-interface** statement at the **[edit system services dhcp-local-server dhcpv6 overrides]**, **[edit system services dhcp-local-server dhcpv6 group *group-name* overrides]**, or **[edit system services dhcp-local-server dhcpv6 group *group-name* interface *interface-name* overrides]** hierarchy levels.

For DHCPv6 relay agent, include the **client-negotiation-match incoming-interface** statement at the **[edit forwarding-options dhcp-relay dhcpv6 overrides]**, **[edit forwarding-options dhcp-relay dhcpv6 group *group-name* overrides]**, or **[edit forwarding-options dhcp-relay dhcpv6 group *group-name* interface *interface-name* overrides]** hierarchy levels.

- **DHCP rate adjustment (MX Series)**—Starting in Junos OS Release 15.1R5, you can use DHCP tags to modify the CLI-configured and RADIUS-configured shaping rate values after a subscriber is instantiated. The new values are conveyed in DHCP option 82, suboption 9 discovery packets. Suboption 9 contains the Internet Assigned Numbers Authority (IANA) DSL Forum VSA (vendor ID 3561).

Configure the shaping rate adjustment controls by including the **dhcp-tags** statement at the **[edit class-of-service adjustment-control-profiles *profile-name* application]**

hierarchy level. Specify the desired rate-adjustment algorithm and set a priority for the DHCP Tags application in the adjustment control profile.

- **Configurable grace period for unresponsive RADIUS servers (MX Series)**—Starting in Junos OS Release 15.1R6, you can use the **timeout-grace** statement at the **[edit access radius-options]** hierarchy level to configure a grace period that determines when an unresponsive RADIUS authentication server is marked as down or unreachable. When the server fails to respond to any of the attempts made for an authentication request, it times out, the time is noted, and the grace period begins. If the server is unresponsive for subsequent authentication requests, the grace period is checked each time the server times out. When the check determines that the grace period has expired, the server is marked as down or unreachable.

You can configure the grace period in the range 0 through 30 seconds; the default is 10 seconds. Use a short grace period to declare servers unavailable sooner and direct requests to available servers. Use a long grace period to give unresponsive servers more opportunities to respond.

In earlier releases, the grace period is 10 seconds and is not configurable.

- **Traffic shaping and tunnel switches (MX Series)**—Starting in Junos OS Release 15.1R6, when a dynamic profile attaches a statically configured firewall filter to an L2TP tunnel switch (LTS) session, the filter polices traffic from the LTS (acting as a LAC) to the ultimate endpoint LNS, in addition to the previously supported traffic from the LAC to the LTS (acting as an LNS). In previous releases, the firewall filter applied to only the traffic from the LAC to the LTS.
- **Enhancement to Gx-Plus Application (MX Series)**—Starting in Junos OS Release 15.1R6, the following enhancements to the Gx-Plus client application on the BNG are available:
 - When a monitored service is deactivated separate from a subscriber logout, the CCR-U indicates that the service is no longer active and includes the service's usage data.
 - The router updates the monitoring key and threshold values when they are received in a RAR message from the PCRF.
 - A CCR-U is sent to the PCRF after the router sends an RAA message in response to an RAR message that requests service activations or deactivations.
 - When the PCRF returns threshold values that are lower than the current values, the new threshold becomes the sum of the current value and the returned value.
 - The PCEF has default minimum threshold values. If the change between the current value and the value returned by the PCRF is less than the minimum value, then the new value is adjusted to the minimum.
 - The CCR-I message includes the Diameter AVP Subscription-Id attribute (443) with the Subscription-Id-Type Diameter AVP sub-attribute (450) set to 4 (END_USER_PRIVATE) and the Subscription-Id-Data Diameter AVP sub-attribute (444) set to **reserved**.

System Logging

- **System log messages to indicate checksum errors on the DDR3 interface**—Starting in Junos OS Release 13.3 R9, two new system log messages, XMCHIP_CMERROR_DDRIF_INT_REG_CHKSUM_ERR_MINOR and XMCHIP_CMERROR_DDRIF_INT_REG_CHKSUM_ERR_MAJOR, are added to indicate memory-related problems on the interfaces to the double data rate type 3 (DDR3) memory. These error messages indicate that an FPC has detected a checksum error, which is causing packet drops.

The following error threshold values classify the error as a major error or a minor error:

- Minor error— 6-254 errors per second
- Major error—255 and more errors per second

User Interface and Configuration

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

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

VPNs

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

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

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

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



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

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

- **Enhanced Group VPNv2 member features (MX10, MX20, MX40, MX80, MX240, MX480, MX960)**—Starting in Junos OS Release 15.1, Group VPNv2 member features have been enhanced to include the following:
 - Accept group domain of interpretation (GDOI) push messages from Cisco group controller/key server (GC/KS) as per RFC 6407.
 - Support for group associated policy (GAP) payload, including activation time delay (ATD) and deactivation time delay (DTD), in push messages from Cisco GC/KS as per RFC 6407.
 - Support standardized push ACK messages from MX Series group member router to Cisco GC/KS as per IETF RFC 8263 [RFC 8263](#).
 - IP Delivery Delayed Detection Protocol. Time-based anti-replay protection for Group VPNv2 data traffic on MX Series group member routers as per IETF draft RFC <http://tools.ietf.org/html/draft-weis-delay-detection-00>.
 - Support for SHA-256 HMAC algorithm for authentication.
 - Support partial fail open for business-critical traffic.
 - Support for control-plane debug traces per member IP address and server IP address.
 - Same gateway for multiple groups, wherein the same local and remote address pair is used for multiple groups.

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

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

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

- **EVPN with VXLAN data plane encapsulation (MX Series)**—Starting in Junos OS Release 15.1R3, MX Series routers can use EVPN with VXLAN encapsulation to provide Layer 2 connectivity for end stations within a Virtualized Network (VN) created by the Contrail virtualization software. The end stations consist of virtual hosts connected to the virtualized server, and non-virtualized bare metal servers connected to top-of-rack platforms. MX Series routers also function as default gateways for the inter-VN traffic among end stations that belong to different VNs. EVPN is used as a Layer 2 overlay solution to provide Layer 2 connections over the IP underlay for the endpoints within a VN whenever Layer 2 connectivity is required by an end station.
- **MVPN source-active upstream multicast hop selection and redundant source improvements**—Starting in Junos OS Release 15.1R3, you can use new configuration statements available at the `[edit protocols mvpn]` hierarchy level to influence the source-active upstream multicast hop selection process. You can use the **umh-selection-additional-input** statement to influence the upstream multicast hop selection by making the MVPN consider some combination of route preference and RSVP tunnel status. You can use the **source-redundancy** statement so that the MVPN acts on all redundant sources sending to a specific group address as the same source.

- See Also**
- [Changes in Behavior and Syntax on page 131](#)
 - [Known Behavior on page 165](#)
 - [Known Issues on page 170](#)
 - [Resolved Issues on page 192](#)
 - [Documentation Updates on page 312](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 320](#)
 - [Product Compatibility on page 329](#)

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 15.1R6 for the M Series, MX Series, and T Series.

- [Authentication, Authorization and Accounting on page 132](#)
- [Class of Service \(CoS\) on page 132](#)
- [General Routing on page 133](#)
- [High Availability \(HA\) and Resiliency on page 133](#)
- [IPv6 on page 135](#)
- [Junos OS XML API and Scripting on page 135](#)
- [Layer 2 Features on page 136](#)
- [Layer 2 VPNs on page 136](#)
- [Management on page 136](#)

- [MPLS on page 136](#)
- [Multicast on page 137](#)
- [Network Management and Monitoring on page 137](#)
- [Platform and Infrastructure on page 138](#)
- [Routing Policy and Firewall Filters on page 138](#)
- [Routing Protocols on page 139](#)
- [Security on page 142](#)
- [Services Applications on page 144](#)
- [Subscriber Management and Services \(MX Series\) on page 146](#)
- [System Logging on page 157](#)
- [System Management on page 164](#)
- [User Interface and Configuration on page 164](#)
- [Virtual Chassis on page 165](#)
- [VLAN Infrastructure on page 165](#)
- [VPNs on page 165](#)

Authentication, Authorization and Accounting

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

Class of Service (CoS)

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

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

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

- **CLI commit check not performed for guaranteed-rate burst size (MX Series)**—Starting in Junos OS Release 15.1R1, the CLI no longer performs a commit check to determine whether the statically configured guaranteed-rate burst size exceeds the shaping-rate burst size. A system log is generated when the guaranteed-rate burst size is higher, whether it is configured statically, dynamically with predefined variables, or by means of a change of authorization request. In earlier releases, a CLI commit check prevents

a static configuration from being used; no checks are performed for the other configuration methods.

General Routing

- **commit synchronize statement is not allowed in batch mode**—When you attempt to execute **commit atomic** in configure batch mode, a warning message is displayed: **warning: graceful-switchover is enabled, commit synchronize should be used**. This is because commit synchronize is not allowed to be given in configure batch mode. In this case, issue the **set system commit synchronize** command followed by **commit**.
- **Modified output of the clear services sessions | display xml command (MX Series)**—In Junos OS Release 14.1X55-D30, the output of the **clear services sessions | display xml** command is modified to include the **<sess-marked-for-deletion>** tag instead of the **<sess-removed>** tag. In releases before Junos OS Release 14.1X55-D30, the output of this command includes the **<sess-removed>** tag. The replacement of the **<sess-removed>** tag with the **<sess-marked-for-deletion>** tag aims at establishing consistency with the output of the **clear services sessions** command that includes the field **Sessions marked for deletion**.
- The **as-path-ignore** command is supported for routing instances starting with Junos OS Release 14.1R8, 14.2R7, and 15.1R4.

High Availability (HA) and Resiliency

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

[See [Configuring a Logical Interface to Be Tracked for a VRRP Group](#) or [Configuring a Route to Be Tracked for a VRRP Group](#).]
- **A check option is added for command request chassis routing-engine master**—Starting in Junos OS Release 15.1, a **check** option available with the **switch**, **release**, and **acquire** options checks the GRES status of the standby Routing Engine before toggling mastership. The **force** option is also removed.

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

[See [show system switchover](#).]
- **Improved command output for determining GRES readiness in an MX Series Virtual Chassis (MX Series routers with MPCs)**—Starting in Junos OS Release 15.1, the **request virtual-chassis routing-engine master switch check** command displays the following

output when the member routers in a Virtual Chassis are ready to perform a graceful Routing Engine switchover (GRES):

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

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

```
Switchover Ready
```

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

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

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



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

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

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

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

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

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

IPv6

- **IPv6 addresses with padded zeros in MIC or MS-MPC system log messages (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1R2, all system log messages originating from MIC or MS-MPC line cards displays padded zeros in IPv6 addresses to make them compatible with MS-DPC line cards. Earlier, the system log messages from MIC or MS-MPC line cards displayed IPv6 addresses with '::' instead of padded zeros.

Junos OS XML API and Scripting

- **Escaping of special XML characters required for request_login (M Series, MX Series, and T Series)**—Beginning with Junos OS Release 15.1R2, you must escape any special characters in the username and password elements of a **request_login** XML RPC request. The following five symbols are considered special characters: greater than (>), less than (<), single quote ('), double quote ("), and ampersand (&). Both entity references and character references are acceptable escape sequence formats. For example, **&** and **&** are valid representations of an ampersand. Previously no escaping of these characters was required.
- **XML output change for show subscribers summary port command (MX Series)**—Starting in Junos OS Release 15.1R5, the display format has changed for the **show subscribers summary port** command to make parsing the output easier. The output is now displayed as in the following example:

```
user@host> show subscribers summary port | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/16.1R2/junos">
  <subscribers-summary-information
    xmlns="http://xml.juniper.net/junos/16.1R2/junos-subscribers">
    <counters junos:style="port-summary">
      <port-name>ge-1/2/0</port-name>
      <port-count>1</port-count>
    </counters>
    <counters junos:style="port-summary">
      <port-name>ge-1/2/1</port-name>
      <port-count>1</port-count>
    </counters>
  </subscribers-summary-information>
</rpc-reply>
```

In earlier releases, that output is displayed as in the following example:

```
user@host> show subscribers summary port | display xml
<rpc-reply xmlns:junos="http://xml.juniper.net/junos/16.1R2/junos">
  <subscribers-summary-information
    xmlns="http://xml.juniper.net/junos/16.1R2/junos-subscribers">
    <counters junos:style="port-summary">
      <port-name>ge-1/2/0</port-name>
      <port-count>1</port-count>
      <port-name>ge-1/2/1</port-name>
      <port-count>1</port-count>
    </counters>
  </subscribers-summary-information>
</rpc-reply>
```

Layer 2 Features

- **Support for configuring MAC move parameters globally (MX Series)**—Starting in Junos OS Release 15.1R4, you can configure parameters for media access control (MAC) address move reporting by including the **global-mac-move** statement and its substatements at the **[edit protocols l2-learning]** hierarchy level. When a MAC address appears on a different physical interface or within a different unit of the same physical interface and this behavior occurs frequently, it is considered a MAC move. You can configure the router to report a MAC address move based on the following parameters: the number of times a MAC address move occurs, a specified period of time over which the MAC address move occurs, and the specified number of times a MAC address move occurs in one second.

Layer 2 VPNs

- **Support for hot standby pseudowire for VPLS instances with LDP (MX Series)**—Starting with Junos OS Release 15.1R2, you can configure a routing device running a VPLS routing instance configured with the Label Distribution Protocol (LDP) to indicate that a hot-standby pseudowire is desired upon arrival of a PW_FWD_STDBY status-tlv. Include the **hot-standby-vc-on** statement at the **[edit routing instances routing-instance-name protocols vpls mesh-group mesh-group-name neighbor address pseudowire-status-tlv]** hierarchy level.
- Starting with Junos OS 15.1R6, the syslog message records a key encryption key (KEK) installation failure when the installation of the KEK security association in a group VPN fails. This is caused by a key server sending an invalid payload. We recommend using the group controller key server (GCKS) on the SRX Series platform as your key server.

Management

- **Support for status deprecated statement in YANG modules (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1R5, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

MPLS

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

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

- **Point-to-multipoint LSP ping echo reply ignored on Juniper side in Cisco-Juniper interoperability (M Series, MX Series, and T Series)**—Currently, in a Juniper-Cisco interoperation network scenario, a point-to-multipoint LSP ping echo reply message from a Cisco device in a different IGP area is dropped on the Juniper device when the

source address of the reply message is an interface address other than the loopback address or router ID.

Starting with Junos OS Release 14.2R6, 15.1R4, 16.1, and later releases, such point-to-multipoint LSP ping echo reply messages are accepted by the Juniper device and the messages get logged as uncorrelated responses.

Multicast

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

Network Management and Monitoring

- **Enhanced service type information in an SNMP MIB walk operation for jnxSpSvcSet**—Starting with releases 13.3R7, 14.1R6, 14.2R4, and 15.1R2, Junos OS provides enhanced service type (SvcType) information in a MIB walk operation for the jnxSpSvcSet MIB table. Stateful firewall, NAT, and IDS service sets are now categorized under the **SFW/NAT/IDS** service type. IPsec services are categorized as **IPSEC** service type, while all other services are grouped as **EXT-PKG**.

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

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

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

- **Change in the output of snmp mib walk of the jnxVpnIfStatus MIB object (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1R4, the `show snmp mib walk jnxVpnIfStatus` command provides information of all interfaces, except the Juniper Networks specific dynamic interfaces.
- **New 64-bit counter of octets for interfaces (M Series, MX Series, and T Series)**—Starting with Release 15.1R3, Junos OS supports two new Juniper Networks

enterprise-specific Interface MIB Extension objects—**ifHCIn1SecOctets** and **ifHCOut1SecOctets**—that act as 64-bit counters of octets passing through an interface.

- **Enhancement for SONET interval counter (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1R3, only the **Current Day Interval Total** output field in the **show interfaces interval** command for SONET interfaces is reset after 24 hours. In addition, the **Previous Day Interval Total** output field displays the last updated time in *hh:mm*.

[See [show interfaces interval](#).]

Platform and Infrastructure

Routing Policy and Firewall Filters

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

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

- **Support for logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol (M Series, MX Series, and T Series)**— Starting with Junos OS Release 15.1R4, you can configure logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol. The queue-depth indicates the number of IP options packets which can be enqueued in the Packet Forwarding Engine logical queue, beyond which it would start dropping the packets.

Routing Protocols

- **Optimization of link-state packets (LSPs) flooding in IS-IS (MX Series)**—Starting in Junos OS Release 15.1R5, flooding of LSPs in IS-IS no longer occurs as a result of the commitment of configuration changes unrelated to IS-IS. Now, when the router is not in the restart state, every time a new LSP is generated after a CLI commit, the contents of the new LSP are compared to the contents of the existing LSP already installed in the link-state database (LSDB) between Intermediate Systems. When the contents of the two LSPs do not match, the system does not process the new LSP or install it in the LSDB, and consequently does not flood it through the IS-IS network. The new behavior does not affect the rebuilding of LSPs after they refresh in the LSDB. No configuration is required to invoke the new behavior.

In earlier releases, IS-IS generates new LSPs even when the configuration changes are not related to IS-IS. Because the new LSPs are flooded across the network and synchronized in the LSDB, this flooding process is time-consuming and CPU intensive in a scaled network environment.

- **Enable forwarding IPv6 solicited router advertisements as unicast**—Beginning with Junos OS Release 15.1, you can configure devices to send router advertisements as unicast in response to the router solicitation message sent by IPv6 routers. In earlier Junos OS releases, IPv6 router advertisements were sent as periodic multicast, which caused a battery drain in all the other devices. A new configuration statement **solicit-router-advertisement-unicast** is introduced at the **[edit protocols router-advertisement interface *interface-name*]** hierarchy level.

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

- **DSCP bit not copied into IPv6 ICMP reply packets (MX Series)**—Beginning with Junos OS Release 15.1, the Differentiated Services code point (DSCP) field from the IPv6 header of the incoming ICMP request packet is copied into the ICMP reply packet. The value of the DSCP field represents the class of service, and transmission of packets is prioritized based on this value. In earlier Junos OS releases, the value of the DSCP field was set to 0, which is undesirable because the class of service information is lost. Junos OS now retains the value of the DSCP field in the incoming packet and copies it into the ICMP reply packet.
- **New option to remove peer loop check (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, a new option **no-peer-loop-check** to remove the peer loop check for private AS numbers is available under the **remove-private** statement at the following hierarchy levels:

```
[edit logical-systems logical-system-name protocols bgp]
[edit protocols bgp]
[edit routing-instances routing-instance-name protocols bgp]
```

- **BGP hides a route received with a label block size greater than 256 (M Series, MX Series, and T Series)**—Starting in Junos OS Release 15.1, when a BGP peer (running Junos OS) sends a route with a label block size greater than 256, the local speaker hides the route and does not re-advertise this route. The output of the **show route detail/extensive hidden/all** displays the hidden route and states the reason as **label block size exceeds max supported value**. In earlier Junos OS releases, when a peer sent

a route with a label block size greater than 256, the routing protocol process (rpd) terminated abnormally.

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

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

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

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

- **Configure and establish targeted sessions with third-party controllers using LDP targeted neighbor (M Series and MX Series)**—Starting with Junos OS Release 15.1, you can configure LDP targeted neighbor to third-party controllers for applications such as route recorder that wants to learn label-FEC bindings of an LSR. LDP targeted neighbor helps to establish a targeted session with controllers for a variety of applications.
- **Enhanced BGP log message when prefix limit is exceeded**—Beginning with Junos OS Release 13.3, BGP generates an enhanced log message when the prefix limit exceeds the configured limit. The log message now includes the instance name in addition to the peer address and address family.

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

- **BGP route is hidden when AS path length is more than the configured maximum AS size**—Beginning with Junos OS Release 13.2, BGP hides a route when the length of the AS path does not match the number of ASs in the route update. In earlier Junos OS releases when a route with AS path size over 2048 was advertised, it could cause session flaps between BGP peers because of the mismatch. Therefore, to avoid session flaps, such routes are now hidden by Junos OS. You can see this behavior when **bgp-error-tolerance** is configured.

If you want BGP to advertise the hidden route to an OSPF neighbor, we recommend to add the AS path statically in the default route configuration. For example:

```
[edit routing-instances instance-name routing options]
user@host# set aggregate route 0.0.0.0/0 as-path path 1267
```

- **BGP link state value modified to 29 (M Series, MX Series, and T Series)**—Starting in Junos OS Release 14.2R3, the value of the BGP **LINK-STATE** (LS) path attribute is modified to 29, which is IANA's officially assigned value. In earlier Junos OS releases, the **LINK-STATE** path attribute had a private value of 99 that was used for interoperability testing with other vendors. The previous versions of BGP LS are not compatible with this new value of BGP LS. Therefore, BGP LS users cannot use unified ISSU with the BGP LS value of 29.
- **New IS-IS adjacency holddown CLI command (MX Series)**—Beginning with Junos OS Release 15.1, a new operational command **show isis adjacency holddown** is introduced to display the adjacency holddown status. This command is useful to verify whether the adjacency holddown is enabled and facilitates troubleshooting when there are adjacency issues due to IS-IS adjacency holddown.
[See [show isis adjacency holddown](#).]
- **Eliminate fe80::/64 direct routes from RIB for IPv6 interfaces**—Beginning with Junos OS Release 15.1, the fe80::/64 direct routes for IPv6 addresses are not installed in the routing table. Therefore, when you issue a **show route** command, the fe80::/64 routes for IPv6 addresses are not displayed in the output. In earlier releases, Junos OS added the fe80::/64 direct routes to the routing table when inet6 family was enabled on an interface. These fe80::/64 direct routes are neither routable nor used for routing decisions and hence their absence in the routing table does not impact any functionality.
- **Support for RFC 5492, Capabilities Advertisement with BGP-4**—Beginning with Junos OS Release 15.1R4, BGP sessions can be established with legacy peers that do not support optional parameters, such as capabilities. In earlier Junos OS releases from 15.1R1 through 15.1R3 and 15.1F1 through 15.1F4, BGP sessions with legacy routers without BGP capabilities was not supported. Starting with Junos OS Release 15.1R4, support for BGP sessions with legacy routers without BGP capabilities is restored.

Security

- **Packet types added for DDoS protection L2TP policers (MX Series routers with MPCs, T4000 routers with FPC5)**—Starting in Junos OS Release 15.1R6, the following eight packet types have been added to the DDoS protection L2TP protocol group to provide flexibility in controlling L2TP packets:

cdn	scccn
hello	sccrq
iccn	stopccn
icrq	unclassified

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

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

- **Changes to distributed denial of service (DDoS) protection protocol groups and packet types (MX Series, T4000 with FPC5)**—Starting in Junos OS Release 15.1, the following syntax changes have been made:
 - The **mlp** protocol group has been modified as follows to provide DDoS protection with full control of the bandwidth:
 - The **aging-exc**, **packets**, and **vxlan** packet types have been removed from the **mlp** protocol group.
 - The **add**, **delete**, and **lookup** packet types have been added to the **mlp** protocol group. These packets correspond to the MAC learning command codes.
 - The **keepalive** protocol group has been renamed to **tunnel-ka**.
 - The **firewall-host** protocol group and the **mcast-copy** packet type in the **unclassified** protocol groups have been removed from the CLI. They are now classified by the internal host-bound classification engine on the line card.
- **Changes to distributed denial of service (DDoS) protection default values for MLP packets (MX Series, T4000 with FPC5)**—Starting in Junos OS Release 15.1, the following default bandwidth (pps) and burst (packets) values apply for MLP packets by line card:

Policer	MPC1, MPC2, MPC5, and MPC6		MPC3, MPC4, and FPC5	
	Bandwidth	Burst	Bandwidth	Burst
aggregate	10,000	20,000	5000	10,000
add	4096	8192	2048	4096

Policer	MPC1, MPC2, MPC5, and MPC6		MPC3, MPC4, and FPC5	
	Bandwidth	Burst	Bandwidth	Burst
delete	4096	8192	2048	4096
lookup	1024	2048	512	1024
unclassified	1024	1024	512	512

- **Changes to distributed denial of service (DDoS) protection flow detection defaults (MX Series, T4000 with FPC5)**—Starting in Junos OS Release 15.1, flow detection defaults to **disabled** for the following protocol groups and packet type, because they do not have typical Ethernet, IP, or IPv6 headers. Global flow detection does not enable flow detection for these groups and the packet type.
 - Protocol groups: **fab-probe**, **frame-relay**, **inline-ka**, **isis**, **jfm**, **mlp**, **pfe-alive**, **pos**, **services**.
 - Packet type: **unclassified** in the **ip-opt** protocol group.
- **Changes to show ddos-protection protocols command output (MX Series, T4000 with FPC5)**—Starting in Junos OS Release 15.1, when you disable DDoS protection policers on the Routing Engine or on an FPC for a specific packet type, an asterisk is displayed next to that field in the CLI output. For example, if you issue the following statements:

```
user@host# set system ddos-protection protocols mlp lookup disable-routing-engine
user@host# set system ddos-protection protocols mlp lookup fpc 1 disable-fpc
```

the fields are marked as in the following sample output:

```
user@host> show ddos-protection protocols mlp lookup
```

```
Currently tracked flows: 0, Total detected flows: 0
* = User configured value
```

```
Protocol Group: MLP
```

```
Packet type: lookup (MLP lookup request)
```

```
Individual policer configuration:
```

```
Bandwidth: 1024 pps
```

```
...
```

```
Routing Engine information:
```

```
Bandwidth: 1024 pps, Burst: 2048 packets, disabled*
```

```
Policer is never violated
```

```
Received: 0 Arrival rate: 0 pps
```

```
Dropped: 0 Max arrival rate: 0 pps
```

```
Dropped by aggregate policer: 0
```

```
FPC slot 1 information:
```

```
Bandwidth: 100% (1024 pps), Burst: 100% (2048 packets), disabled*
```

```
Policer is never violated
```

```
Received: 0 Arrival rate: 0 pps
```

```
Dropped: 0 Max arrival rate: 0 pps
```

```
Dropped by aggregate policer: 0
```

```
Dropped by flow suppression: 0
```

Services Applications

- **Support for configuring TWAMP servers on routing instances (MX Series)**—Starting in Junos OS Release 15.1, you can specify the TWAMP servers on specific routing instances, instead of associating the TWAMP server at the system level. To apply the TWAMP server to a routing instance configured on a router, include the **routing-instance-list instance-name port port-number** statement at the **[edit services rpm twamp server]** hierarchy level. The port number of the specified routing instance is used for TWAMP probes that are received by a TWAMP server. The default routing instance is Internet routing table inet.0. If you do not specify a routing instance, the TWAMP probe applies to all routing instances. To apply the TWAMP probe to only the default routing instance, you must explicitly set the value of instance-name to the default. If an interface is not part of any routing instance, the default port is used for TWAMP probes. You can configure up to 100 routing instances for a TWAMP server.
- **Optional inclusion of Flags field in DTCP LIST messages (MX Series)**—Starting in Junos OS Release 15.1, the Flags field is not a required parameter in the DTCP LIST message. The LIST request is not rejected if the LIST message does not contain the Flags field. If the DTCP LIST message contains the Flags field, the value of that field is processed. If the LIST message does not contain the Flags field, the CRITERIA field parameter is used for the Flags field.
- **Change in support for service options configuration on service PICs at the MS and AMS interface levels (MX Series)**—Starting in Junos OS Release 15.1, when a multiservices PIC (**ms-** interface) is a member interface of an AMS bundle, you can configure the service options to be applied on the interface only at the **ms-** interface level or the AMS bundle level by including the **services-options** statement at the **[edit interfaces interface-name]** hierarchy level at a point in time. You cannot define service options for a service PIC at both the AMS bundle level and at the **ms-** interface level simultaneously. When you define the service options at the MS level or the AMS bundle level, the service options are applied to all the service-sets, on the **ms-** interface or the AMS interface defined at **ms-fpc/pic/port.logical-unit** or **amsN**, respectively.
- **Changes in the format of session open and close system log messages (MX Series routers with MS-MICs and MS-MPCs)**—Starting with Junos OS Release 15.1, with the Junos OS Extension-Provider packages installed and configured on the device for MS-MPCs and MS-MICs, the formats of the MSVCS_LOG_SESSION_OPEN and MSVCS_LOG_SESSION_CLOSE system log messages are modified to toggle the order of the destination IPv4 address and destination port address displayed in the log messages to be consistent and uniform with the formats of the session open and close logs of MS-DPCs.
- **Support for bouncing service sets for dynamic NAT (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 15.1, for service sets associated with aggregated multiservices (AMS) interfaces, you can configure the **enable-change-on-ams-redistribution** statement at the **[edit services service-set service-set-name service-set-options]** hierarchy level to enable the service set to be bounced (reset) for dynamic NAT scenarios (dynamic NAT, NAT64, and NAT44) when a member interface of an AMS bundle rejoins or a member interface failure occurs. When a member interface fails, the application resources (NAT pool in the case of dynamic NAT scenarios) and traffic load need to be rebalanced. For application

resources to be rebalanced, which is the NAT pool for dynamic NAT environments, the NAT pool is split and allocated by the service PIC daemon (spd).

- **Changed range for maximum lifetime for PCP mapping**—Starting in Junos OS Release 15.1, the range for the maximum lifetime, in seconds, for PCP mapping that you can configure by using the `mapping-lifetime-max` statement at the `[edit services pcp]` hierarchy level is modified to be from 0 through 4294667, instead of the previous range from 0 through 2147483647.
- **Change in the test-interval range for RPM tests (MX Series)**—Starting in Junos OS Release 15.1R2, the minimum period for which the RPM client waits between two tests (configured by using the `test-interval` statement at the `[edit services rpm probe owner test test-name]` hierarchy level is modified to be 1 second instead of 0 seconds. Also, if you do not configure the test interval, the default value is 0 seconds. A test interval of 0 seconds causes the RPM test to stop after one iteration.
- **Change to show services nat pool command output**—Starting in Junos OS Release 15.1R3, the `show services nat pool` command output includes this new field: AP-P port limit allocation errors. When AP-P is configured, this field indicates the number of out-of-port errors that are due to a configured limit for the number of allocated ports in the `limit-ports-per-address` statement at the `[edit services nat pool nat-pool-name]` hierarchy level.
- **Class pcp-logs and alg-logs are not configured for ms-interface (MX Series)**—Starting with Junos OS release 15.1R3, for multiservices (ms-) interfaces, you cannot configure system logging for PCP and ALGs by including the `pcp-logs` and `alg-logs` statements at the `[edit services service-set service-set-name syslog host hostname class]` hierarchy level. An error message is displayed if you attempt to commit a configuration that contains the `pcp-logs` and `alg-logs` options to define system logging for PCP and ALGs for ms- interfaces.
- **Support for deterministic NAPT (MX Series)**—You can configure deterministic port block allocation for Network Address Port Translation (NAPT) on MX Series routers with MS-MPCs or MS-MICs. By configuring deterministic NAPT, you ensure that translation of the internal host IP (private IP to public IP and vice versa) is deterministic, thus eliminating the need for address translation logging for each connection. To use deterministic port block allocation, you must specify `deterministic-napt44` as the translation type in your NAT rule.
- **Anycast address 0/0 must not be accepted in the from-clause of Detnat rule (MX Series)**—Starting with Junos OS Release 15.1R4, for multiservices (ms-) interfaces, anycast configuration is not allowed as the source-address when translation type is deterministic NAT.
- **Deprecated security idp statements (MX Series)**—Starting in Junos OS 15.1R6, the `[edit security idp]` configuration statements are deprecated.
- **Changes to the PGCP service (MX, M and T Series)**—Starting in Junos OS Release 15.1R6, the Packet Gateway Control Protocol (PGCP) is removed from the list of processes during boot. The configuration statements, commands, and options for PGCP process are deprecated. In earlier releases, PGCP process configures the PGCP that is required for the border gateway function (BGF) feature.

Subscriber Management and Services (MX Series)

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

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

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

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

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

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

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

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

- **interface-text-description**—Interface description text string
- **stacked-vlan**—Stacked VLAN ID
- **vlan**—VLAN ID

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

- **Unique RADIUS NAS-Port attributes (MX Series)**—Starting in Junos OS Release 15.1, you can configure unique values for the RADIUS NAS-Port attribute (attribute 5), to ensure that a single NAS-Port attribute is not used by multiple subscribers in the network. You can create NAS-Port values that are unique within the router only, or that

are unique across all MX Series routers in the network. To create unique NAS-Port attributes for subscribers, the router uses an internally generated number and an optional unique chassis ID, which you specify. The generated number portion of the NAS-Port provides uniqueness within the router only. The addition of the optional chassis ID configuration ensures that the NAS-Port is unique across all MX Series routers in the network.

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

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

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

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

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

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

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

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

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

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

- **Changes to ANCP triggering of RADIUS immediate interim accounting updates (MX Series)**—Starting in Junos OS Release 15.1, the AAA daemon immediately sends a RADIUS interim-accounting request to the RADIUS server when it receives notification of ANCP actual downstream or upstream data rate changes, even when the **update-interval** statement is not included in the subscriber session access profile. In earlier releases, the **update-interval** statement is required. This feature still requires that the **ancp-speed-change-immediate-update** statement is included in the access profile.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

- **Local DNS configurations available when authentication order is set to none (MX Series)**—Starting in Junos OS Release 15.1R2, subscribers get the DNS server addresses when both of the following are true:
 - The authentication order is set to **none** at the **[edit access profile *profile-name* authentication-order]** hierarchy level.
 - A DNS server address is configured locally in the access profile with the **domain-name-server**, **domain-name-server-inet**, or **domain-name-server-inet6** statement at the **[edit access profile *profile-name*]** hierarchy level.

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

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

clear services l2tp destination	show services l2tp destination extensive
clear services l2tp session	show services l2tp destination statistics
clear services l2tp tunnel	show services l2tp session extensive
	show services l2tp session statistics
	show services l2tp summary statistics
	show services l2tp tunnel extensive
	show services l2tp tunnel statistics



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

- **Improved result code reporting in stopCCN and CDN messages (MX Series)**—Starting in Junos OS Release 15.1R3, the LAC provides more accurate result codes and always includes error messages in the Result-Error Code AVP (1) included in the stopCCN and CDN messages that it sends to the LNS. Packet captures display the relevant information in the **Result code**, **Error code**, and **Error Message** fields of the AVP.

In earlier releases, the result code does not provide sufficient information about the cause of the event and the error message is omitted for some result codes.

- **Including termination reason for user logout events (MX Series)**—Starting in Junos OS Release 15.1R2, when you enable the user-access flag at the **[edit system**

`processes general-authentication-service traceoptions]` hierarchy level, the system log messages generated for `authd` include a termination reason for user logout events. In earlier releases, the log does not report any termination reasons.

Sample output before the behavior change:

```
Aug  2 15:10:28.181293 UserAccess:zf@example.com session-id:19 state:log-out
ge-1/1/0.100:100-1
```

Sample output after the behavior change:

```
Aug  6 21:15:55.106031 UserAccess:zf@example.com session-id:3 state:log-out
ge-1/2/0.1:1 reason: ppp lcp-peer-terminate-term-req
Aug  6 21:16:42.654181 UserAccess:user234@example.com session-id:4 state:log-out
ge-1/2/0.1:1 reason: ppp lower-interface-down
Aug  6 21:17:43.991585 UserAccess:duser9five@example.com session-id:5
state:log-out ge-1/2/0.1:1 reason: aaa shutdown-session-timeout
```

- **Change in displayed value of LCP State field for tunneled subscriber sessions (MX Series)**—Starting in Junos OS Release 15.1R3, when a subscriber session has been tunneled from the LAC to the LNS, the **LCP State** field displayed by the **show interfaces pp0.unit** command has a value of **Stopped**, which correctly reflects the actual state of the LCP negotiation (because at this stage LCP is terminated at the LNS).

In earlier releases, this field incorrectly shows a value of **Opened**, reflecting the state of LCP negotiation before tunneling started. In earlier releases, you must issue the **show ppp interface.unit** command to display the correct LCP state.

- **Change in Routing Engine-based CPCD (MX Series)**—Starting in Junos OS Release 15.1R3, you must specify a URL with the **redirect** statement. You must also specify **destination-address address** with the **rewrite** statement. In earlier releases, you can successfully commit the configuration without these options.
- **Increased maximum limits for accounting and authentication retries and timeouts (MX Series)**—Starting in Junos OS Release 15.1R3, you can configure a maximum of 100 retry attempts for RADIUS accounting (**accounting-retry** statement) or authentication (**retry** statement). In earlier releases, the maximum value is 30 retries. You can also configure a maximum timeout of 1000 seconds for RADIUS accounting (**accounting-timeout** statement) or authentication (**timeout** statement). In earlier releases the maximum timeout is 90 seconds.



NOTE: The maximum retry duration (the number of retries times the length of the timeout) cannot exceed 2700 seconds. An error message is displayed if you configure a longer duration.

- **Support for longer CHAP challenge local names (MX Series)**—Starting in Junos OS Release 15.1R3, the supported length of the CHAP local name is increased to 32 characters. In earlier releases, only 8 characters are supported even though the CLI allows you to enter a longer name. You can configure the name with the **local-name** statement at the **[edit dynamic-profiles profile-name interfaces pp0 unit**

"\$junos-interface-unit" ppp-options] or **[edit dynamic-profiles profile-name interfaces "\$junos-interface-ifd-name" unit "\$junos-interface-unit" ppp-options]** hierarchy levels. The maximum length of the local name for PAP authentication remains unchanged at 8 characters.

- **Change to test aaa commands (MX Series)**—Starting in Junos OS Release 15.1R4, the following changes have been made to the **test aaa ppp user**, **test aaa dhcp user**, and **test aaa authd-lite user** commands:
 - Attributes not supported by Junos OS no longer appear in the output.
 - The Virtual Router Name and Routing Instance fields have been combined into the new Virtual Router Name (LS:RI) field. The value of this field matches the Juniper Networks Virtual-Router VSA (26-1), if present; otherwise the field displays **default:default**.
 - The value for any attribute that is not received (except for 26-1), or set locally, is displayed as **<not set>**.
 - The Redirect VR Name field has been renamed to Redirect VR Name (LS:RI).
 - In the CLI output header section, the Attributes area has been renamed to User Attributes.
 - Supported attributes now always appear in the display, even when their values are not set.
 - The IGMP field has been renamed to IGMP Enable.
 - The IGMP Immediate Leave and the MLD Immediate Leave default values have changed from **disabled** to **<not set>**.
 - The Chargeable user identity value has changed from an integer to a string.
 - The Virtual Router Name field has been added to the display for the DHCP client.
- **Change to using the UID as part of a variable expression (MX Series)**—Starting in Junos OS Release 15.1R4, you cannot use the UID (the unique identifier of variables defined in dynamic profiles) as part of a variable expression, because the hierarchy of evaluation is as follows:
 - The user variable expressions are first evaluated for the UIDs to be resolved.
 - If the expression contains UIDs, it might result in unpredictable results.

Using a variable expression with a UID now results in a commit check failure.

- **Subscriber management 64-bit mode support (MX Series)**—Starting in Junos OS Release 15.1R4, subscriber management is now supported when the routing protocol daemon (rpd) is running in 64-bit mode. In earlier releases, subscriber management support required rpd to run in 32-bit mode.
- **Subscriber secure policies and service change of authorization requests (MX Series)**—Starting in Junos OS Release 15.1R4, a subscriber secure policy cannot be instantiated by a CoA that includes any other subscriber service activation or deactivation. Use a separate CoA to apply a subscriber secure policy.

- **Configuration support for L2TP hashing (MX Series)**—Starting in Junos OS Release 15.1R4, you can enable or disable the inclusion of the L2TP tunnel ID and session ID in the L2TP packet header in the hash computation for L2TP data packets on an aggregated Ethernet interface to more accurately balance the traffic load over multiple active links. By default, tunnel and session IDs are not considered. To enable the IDs to be used, include the `l2tp-tunnel-session-identifier` statement at the `[edit forwarding-options enhanced-hash-key family inet]` hierarchy level. To disable the inclusion of the IDs, remove the statement from your configuration.

In earlier releases, tunnel and session IDs are included by default for L2TP hashing over aggregated Ethernet links and cannot be disabled.

- **Extended range for RADIUS request rate (MX Series)**—Starting in Junos OS Release 15.1R4, the range for the `request-rate` statement at the `[edit access radius-options]` hierarchy level has been extended to 100 through 4000 requests per second. In earlier releases, the range is 500 through 4000 requests per second. The default value is unchanged at 500 requests per second.
- **VLAN demux interfaces over pseudowire interfaces (MX Series)**—Starting in Junos OS Release 15.1R3, VLAN demux interfaces are supported over pseudowire subscriber logical interfaces.
- **Error messages generated for L2TP access concentrator (LAC) logins can be prevented from appearing in the syslogs**—Starting with Junos OS Release 15.1R4, setting the syslogs log level to WARNING or higher prevents error messages generated for Layer 2 Tunneling Protocol (L2TP) subscribers from appearing in the syslogs. The syslogs are L2TP packet statistics counters (Rx/Tx) that are displayed every minute. If no packets are received or L2TP is not configured, these messages do not appear in the syslogs.

In earlier releases, the severity of the log level was ERROR, which now has changed to NOTICE. The error messages are filtered out if the log level is set to WARNING or higher (ERROR, CRITICAL, ALERT, or EMERGENCY). Setting the log level to NOTICE or lower (INFORMATIONAL or DEBUG) allows the error messages to appear in the syslogs.

- **Configuring a pseudowire subscriber interface for a logical tunnel (MX Series)**—Starting in Junos OS release 15.1R4, you can configure a pseudowire subscriber interface and anchor it to a logical tunnel interface without explicitly specifying the tunnel bandwidth. In earlier releases, if you do not explicitly specify the tunnel bandwidth, or the tunnel bandwidth is anything other than 1G or 10G, the pseudowire interface is not created.
- **L2TP statistics now included in the output of the show system subscriber-management statistics command**—Starting in Junos OS Release 15.1R4, a new option displays the L2TP plugin statistics in the output of the `show system subscriber-management statistics` command.

The possible completions for the `show system subscriber-management statistics` command are:

- `<[Enter]>` executes this command
- `all`—Displays all statistics

- **dhcp**—Displays the DHCP statistics
- **dvlan**—Displays the DVLAN statistics
- **l2tp**—Displays the L2TP statistics
- **ppp**—Displays the PPP statistics
- **pppoe**—Displays the PPPoE statistics
- **/**—Pipes through a command
- **Changes to the test aaa ppp user command (MX Series)**—Starting in Junos OS Release 15.1, the following changes have been made to the **test aaa ppp user** command:
 - Subscriber management supports only the default logical system.
 - Two contexts that now need to be considered:
 - AAA context:
 - The context (LS:RI) is used to authenticate the subscriber.
The Virtual Router Name and the Routing Instance attributes have been combined into a single attribute in the (LS:RI) notation.
 - The **test aaa ppp** command specified on the command line has the following possible completions:
 - **agent-remote-id**—Tests the DSL Forum Agent Remote Id (VSA 26-2)
 - **l2tp-terminate-code**—Tests the L2TP terminate code associated subscriber termination
 - **logical-system**—Tests the logical system in which the user is authenticated
 - **password**—Tests the password associated with the username
 - **profile**—Tests the access profile name associated with the user
 - **routing-instance**—Tests the routing instance in which the user is authenticated
 - **service-type**—Tests the Service type (1-255)
 - **terminate-code**—Tests the PPP terminate code associated with subscriber termination
 - **user**—Tests the username
 - Subscriber context:
 - The context (LS:RI) in which the subscriber is placed. This is established by either Juniper Networks VSA Virtual-Router (26-1) or Juniper Networks VSA Redirect-VRouter-Name (26-25) using (LS:RI) notation, where the routing instance may be different than the AAA context routing instance.
 - Both contexts perform subscriber placement, but the redirect re-authenticates with the RADIUS server in the subscriber context (for example, for L3 wholesale) and may be used for duplicate accounting.
 - Changed items:

- The Chargeable user identity value has changed from int to string.
 - All **not set**, **NULL**, and **Null** outputs have been changed to **<not set>**.
 - Almost all display attributes now show **<not set>** when no value exists and zero is not a valid value for those attributes.
 - Both of the IGMP_Immediate_Leave and MLD Immediate Leave default values have changed from **disabled** to **<not set>**.
 - The Redirect VR Name display format for PPP clients has been changed to (LS:RI) notation.
 - The Virtual Router Name display format for PPP clients has been changed to (LS:RI) notation.
 - Added items:
 - Virtual Router Name has been added to the display for the DHCP client.
 - Removed items:
 - The Routing Instance display has been removed from the output.
 - The Ignore_DF_Bit display has been removed from the output.
 - Both Ingress Statistics and Egress Statistics have been removed from the output.
 - Renamed items:
 - The IGMP display has been renamed to IGMP Enable.
 - Attributes has been renamed User Attributes.
 - **RADIUS VSA support for IANA Private Enterprise Number 311 primary and secondary DNS servers (MX Series)**—Starting in Junos OS Release 15.1, the Junos OS AAA implementation supports RADIUS VSAs that identify the primary and secondary DNS servers for IANA private enterprise number 311 (Microsoft Corporation). The two VSAs are shown in the following list, and are described in RFC 2548, *Microsoft Vendor-specific RADIUS Attributes*:
 - MS-Primary-DNS-Server (VSA 26-28)—The 4-octet address of the primary Domain Name Server. This VSA can be included in Access-Accept and Accounting-Request packets.
 - MS-Secondary-DNS-Server (VSA 26-29)—The 4-octet address of the secondary Domain Name Server. This VSA can be included in Access-Accept and Accounting-Request packets.
- [See [RADIUS Support for Microsoft Corporation VSAs for DNS Server Addresses.](#)]
- **Support deprecated for retaining DHCP subscriber binding during interface deletion (MX Series)**—Starting in Junos OS Release 15.1R4, when enhanced subscriber management is enabled, the MX Series routers no longer support the retention of DHCP bindings during an interface deletion. The **maintain-subscriber** stanza at the **[edit system services subscriber-management]** hierarchy level is deprecated for MX Series routers.

- **Automatic limit set for transmit window size (MX Series)**—Starting in Junos OS Release 15.1R5, when the LAC receives a receive window size of more than 128 in the Start-Control-Connection-Reply (SCCRP) message, it sets the transmit window size to 128 and logs an Error level syslog message.

In earlier releases, the LAC accepts any value sent in the Receive Window Size attribute-value pair (AVP 10) from an L2TP peer. Some implementations send a receive window size as large as 65530. Accepting such a large value causes issues in the L2TP congestion/flow control and slow start. The router may run out of buffers because it can support only up to a maximum of 60,000 tunnels.

- **Change in PPP keepalive interval for inline services subscribers (MX Series)**—Starting in Junos OS Release 15.1R5, you can configure the PPP keepalive interval for subscriber services in the range 1 second through 600 seconds. Subscriber PPP keepalives are handled by the Packet Forwarding Engine. If you configure a value greater than 600 seconds, the number is accepted by the CLI, but the Packet Forwarding Engine limits the interval to 600 seconds. The interval is configured in a PPP dynamic profile with the **interval** statement at the **[edit dynamic-profiles profile-name interfaces pp0 unit \$junos-interface-unit keepalives]** hierarchy level.

In earlier Junos OS releases, the range is from 1 second through 60 seconds. The Packet Forwarding Engine limits any higher configured value to an interval of 60 seconds.

PPP keepalives for non-subscriber services are handled by the Routing Engine with an interval range from 1 second through 32,767 seconds.

- **DNS servers displayed by the show subscribers extensive command (MX Series)**—Starting in Junos OS Release 15.1R6, the display of DHCP Domain Name System (DNS) by the **show subscribers extensive** command has changed. When DNS addresses are configured at multiple levels, the command displays only the preferred address according to this order of precedence: RADIUS > access profile > global access. The command does not display DNS addresses configured as DHCP local pool attributes.

DNS addresses from RADIUS appear in the following fields: Primary DNS Address, Secondary DNS Address, IPv6 Primary DNS Address, and IPv6 Secondary DNS Address.

DNS addresses from the access profile or the global access configuration appear in the following fields: Domain name server inet and Domain name server inet6.

In earlier releases, the command displays only DHCP DNS addresses provided by RADIUS.

- **Default L2TP resynchronization method changed and statement deprecated (MX Series)**—Starting in Junos OS Release 15.1R6, the default resynchronization method for L2TP peers in the event of a control connection failure is changed to silent failover. In earlier releases, the default method is failover-protocol-fall-back-to-silent-failover. The silent failover method is preferred because it does not keep tunnels open without traffic flow, waiting for the failed peer to recover and resynchronize. You can use the new **failover-resync** statement at the **edit services l2tp tunnel** hierarchy level to specify either failover protocol or silent failover as the resynchronization method.

Because silent failover is now the default, the **disable-failover-protocol** statement is no longer needed and has been deprecated. If you upgrade to this release with a

configuration that includes this statement, it is supported, but the CLI notifies you it is deprecated.

- **IPv6 Link Local Addresses Assigned to Underlying Static Demux Interfaces (MX Series)**—Starting in Junos OS Release 15.1R6, when you are using Router Advertisement for IPv6 subscribers on dynamic demux interfaces that run over underlying static demux interfaces, configure the software to use the same link-local address for both interfaces. In this case, the link-local address for the underlying interface should be based the MAC address of the underlying interface. The following statement causes the system to assign an address using the 64-bit Extended Unique Identifier (EUI-64) as described in RFC 2373:

```
system {
  demux-options {
    use-underlying-interface-mac
  }
}
```

- **Traffic shaping and tunnel switches (MX Series)**—Starting in Junos OS Release 15.1R6, when a dynamic profile attaches a statically configured firewall filter to an L2TP tunnel switch (LTS) session, the filter polices traffic from the LTS (acting as a LAC) to the ultimate endpoint LNS, in addition to the previously supported traffic from the LAC to the LTS (acting as an LNS). In previous releases, the firewall filter applied to only the traffic from the LAC to the LTS.

System Logging

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

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

- **New JSERVICES system log messages (MX Series)**—In Junos OS Release 15.1 R3, you can configure MX Series routers with MS-MPCs to log the following messages:

Table 2: JSERVICES System Logs

Name	System Log Message	Description	Severity
------	--------------------	-------------	----------

Table 2: JSERVICES System Logs (continued)

JSERVICES_ALG_FTP_ACTIVE_ACCEPT	software-string <i>src-ip:src-port</i> [<i>xlated-src-ip:xlated-src-port</i>]->[<i>xlated-dst-ip:</i> <i>xlated-dst-port</i>] <i>dst-ip:dst-port (protocol-name)</i>	A FTP data connection from client to server is established. The matching packet contains the indicated information about its protocol name, application, source (logical interface name, IP address, and port number), and destination (IP address and port number). If the flow requires Network Address Translation (NAT) services, NAT information appears in the message.	LOG_NOTICE
JSERVICES_ALG_FTP_PASSIVE_ACCEPT	software-string <i>src-ip:src-port</i> [<i>xlated-src-ip:xlated-src-port</i>]->[<i>xlated-dst-ip:</i> <i>xlated-dst-port</i>] <i>dst-ip:dst-port (protocol-name)</i>	A FTP data connection from server to client is established. The matching packet contains the indicated information about its protocol name, application, source (logical interface name, IP address, and port number), and destination (IP address and port number). If the flow requires Network Address Translation (NAT) services, NAT information appears in the message.	LOG_NOTICE
JSERVICES_DROP_FLOW_DELETE	software-string <i>src-ip:src-port</i> [<i>xlated-src-ip:xlated-src-port</i>]->[<i>xlated-dst-ip:</i> <i>xlated-dst-port</i>] <i>dst-ip:dst-port (protocol-name)</i>	The session with the indicated characteristics is removed and it had drop flow. The NAT data is available in the message if the session requires Network Address Translation (NAT).	LOG_NOTICE

Table 2: JSERVICES System Logs (continued)

JSERVICES_ICMP_ERROR_DROP	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The ICMP error packet was dropped because it did not belong to an existing flow.	LOG_NOTICE
JSERVICES_ICMP_HEADER_LEN_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The ICMP packet was discarded because the length field in the packet header was shorter than the minimum 8 bytes required for an ICMP packet.	LOG_NOTICE
JSERVICES_ICMP_PACKET_ERROR_LENGTH	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The ICMP packet was discarded because the packet contained fewer than 48 bytes or more than 576 bytes of data.	LOG_NOTICE
JSERVICES_IP_FRAG_ASSEMBLY_TIMEOUT	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet and all related IP fragments previously received were discarded because all fragments did not arrive within the reassembly timeout period of four seconds.	LOG_NOTICE
JSERVICES_IP_FRAG_OVERLAP	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the contents of two fragments overlapped.	LOG_NOTICE
JSERVICES_IP_PACKET_CHECKSUM_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because checksum was incorrect.	LOG_NOTICE
JSERVICES_IP_PACKET_DST_BAD	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because its destination address was either a multicast address or was in the range reserved for experimental use (248.0.0.0 through 255.255.255.254).	LOG_NOTICE

Table 2: JSERVICES System Logs (continued)

JSERVICES_IP_PACKET_FRAG_LEN_INV	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because the length of a fragment was invalid.	LOG_NOTICE
JSERVICES_IP_PACKET_INCORRECT_LEN	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The IP packet is discarded because packet length was invalid.	LOG_NOTICE
JSERVICES_IP_PACKET_LAND_ATTACK	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its source and destination address for the packet were the same (referred as land attack).	LOG_NOTICE
JSERVICES_IP_PACKET_LAND_PORT_ATTACK	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its source and destination address for the packet were the same and also its source and destination ports were same (referred as land port attack).	LOG_NOTICE
JSERVICES_IP_PACKET_NOT_VERSION_4	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet version was not IP version 4 (IPv4).	LOG_NOTICE
JSERVICES_IP_PACKET_NOT_VERSION_6	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet version was not IP version 6 (IPv6).	LOG_NOTICE
JSERVICES_IP_PACKET_PROTOCOL_ERROR	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because it used invalid IP protocol.	LOG_NOTICE
JSERVICES_IP_PACKET_SRC_BAD	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because its source address was one of the following: (1) a multicast address (2) a broadcast address (3) in the range 248.0.0.0 through 255.255.255.254, which is reserved for experimental use.	LOG_NOTICE

Table 2: JSERVICES System Logs (continued)

JSERVICES_IP_PACKET_TTL_ERROR	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet with the indicated characteristics is discarded because the packet had a time-to-live (TTL) value of zero.	LOG_NOTICE
JSERVICES_IP_PACKET_TOO_LONG	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet was discarded because the packet contained more than 64 kilobytes (KB) of data (referred to as a ping-of-death attack).	LOG_NOTICE
JSERVICES_IP_PACKET_TOO_SHORT	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet did not contain the minimum amount of data required.	LOG_NOTICE
JSERVICES_NO_IP_PACKET	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	Packet received was not an IPv4 or IPv6 packet.	LOG_NOTICE
JSERVICES_SYN_DEFENSE	proto <i>protocol-id</i> (<i>protocol-name</i>), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description	The packet with the indicated characteristics was discarded because the Transmission Control Protocol (TCP) handshake that is used to establish a session did not complete within the set time limit. The time limit is set by the 'open-timeout' statement at the [edit interfaces <services-interface> services-options] hierarchy level. If the time limit is not set, the session uses the default timeout value.	LOG_NOTICE

Table 2: JSERVICES System Logs (continued)

JSERVICES_SFW_NO_POLICY	<i>source-ip:destination-ip</i> No policy	The stateful firewall received packets with the indicated source and destination addresses. There was no matching policy for the traffic.	LOG_NOTICE
JSERVICES_SFW_NO_RULE_DROP	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The stateful firewall discarded the packet with the indicated characteristics, because the packet did not match any stateful firewall rules. In this case, the default action is to discard the packet. The discarded packet contained the indicated information about its protocol (numerical identifier and name), source (logical interface name, IP address, and port number), and destination (IP address and port number).	LOG_NOTICE
JSERVICES_TCP_FLAGS_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the flags in the packet were set in one of the following combinations: (1) FIN and RST (2) SYN and one or more of FIN, RST, and URG.	LOG_NOTICE
JSERVICES_TCP_HEADER_LEN_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the length field in the packet header was shorter than the minimum 20 bytes required for a TCP packet.	LOG_NOTICE
JSERVICES_TCP_NON_SYN_FIRST_PACKET	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The TCP packet was discarded because it was the first packet in the TCP session but the SYN flag was not set.	LOG_NOTICE

Table 2: JSERVICES System Logs (continued)

JSERVICES_TCP_PORT_ZERO	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the source or destination port specified in the packet was zero.	LOG_NOTICE
JSERVICES_TCP_SEQNUM_AND_FLAGS_ZERO	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the packet's sequence number was zero and no flags were set.	LOG_NOTICE
JSERVICES_TCP_SEQNUM_ZERO_FLAGS_SET	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The packet was discarded because the packet's sequence number was zero and one or more of the FIN, PSH, and URG flags were set.	LOG_NOTICE
JSERVICES_UDP_HEADER_LEN_ERROR	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The UDP packet was discarded because the length field in the packet header was shorter than the minimum 8 bytes required for an UDP packet.	LOG_NOTICE
JSERVICES_UDP_PORT_ZERO	<i>proto protocol-id (protocol-name), source-interface-name:source-address:source-port -> destination-address:destination-port, event-description</i>	The UDP packet was discarded as the source or destination port specified in the packet was zero.	LOG_NOTICE

System Management

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

User Interface and Configuration

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

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

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

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

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

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

- **New warning message for the configurational changes to extend-size (M Series, MX Series, and T Series)**—Starting with Junos OS Release 15.1R2, any operation on the **system configuration-database extend-size** configuration statement, such as **deactivate**, **delete**, or **set**, generates the following warning message:

Change in 'system configuration-database extend-size' will be effective at next reboot only.

Virtual Chassis

- **SNMP MIB walk on MX series Virtual Chassis**—Starting with Junos OS Release 15.1R3, `snmp mib walk` operations no longer return invalid PCMCIA card information for Routing Engines on MX Series Virtual Chassis.

VLAN Infrastructure

- **ACI and ARI from PADI messages included in Access-Request messages for VLAN authentication (MX Series)**—Starting in Junos OS Release 15.1R5, when the PPPoE PADI message includes the agent circuit identifier (ACI), agent remote identifier (ARI), or both, these attributes are stored in the VLAN shared database entry. If the VLAN needs to be authenticated, then these attributes are included in the RADIUS Access-Request message as DSL Forum VSAs 26-1 and 26-2, respectively (vendor ID 3561). The presence of these attributes in the Access-Request enables the RADIUS server to act based on the attributes.

VPNs

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

- See Also**
- [New and Changed Features on page 80](#)
 - [Known Behavior on page 165](#)
 - [Known Issues on page 170](#)
 - [Resolved Issues on page 192](#)
 - [Documentation Updates on page 312](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 320](#)
 - [Product Compatibility on page 329](#)

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R6 for the M Series, MX Series, and T Series.

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

- [Hardware on page 166](#)
- [Forwarding and Sampling on page 167](#)

- [Interfaces and Chassis on page 167](#)
- [MPLS on page 167](#)
- [Network Management and Monitoring on page 167](#)
- [Subscriber Management and Services \(MX Series\) on page 167](#)
- [System Logging on page 169](#)
- [VPNs on page 170](#)

Hardware

- **Support for MIC-3D-8OC3-2OC12-ATM Revision 22 and later**—MIC-3D-8OC3-2OC12-ATM Revision 22 is supported only by the following Junos OS releases:

- Junos OS Release 12.3—12.3R9 and later
- Junos OS Release 13.3—13.3R6 and later
- Junos OS Release 14.1—14.1R4 and later
- Junos OS Release 14.2—14.2R3 and later
- Junos OS Release 15.1 and later

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

- **The options `alarm low-light-alarm` and `warning low-light-warning` might not work (MX Series)**—The `alarm low-light-alarm` and `warning low-light-warning` options at the `[edit interfaces interface-name optics-options]` hierarchy level might not work for the 10-Gigabit Ethernet and 100-Gigabit Ethernet interfaces of MPC3, MPC4, MPC5, MPC6, MPC7E, MPC8E, and MPC9E on MX Series 5G Universal Routing Platforms. These options might not work on MPC3E-3D-NG, MPC3E-3D-NG-Q, MPC2E-3D-NG, and MPC2E-3D-NG-Q if they are installed with Junos Continuity software.

This is a known behavior and has no impact on the performance of these line cards.

Forwarding and Sampling

- **Improved Packet Forward Engine performance (MX Series)**--Starting in Junos OS Release 15.1R5,, a new mechanism is added to the Packet Forwarding Engine to improve forwarding performance. A noticeable behavior of the mechanism is to increase the CPU utilization periodically.

Interfaces and Chassis

- **Reordering of MAC addresses after a Routing Engine switchover**--In Junos OS Release 14.2 and later, if you configure multiple aggregated Ethernet interfaces, the MAC address of the aggregated Ethernet interfaces displayed in the **show interfaces ae number** command output might get reordered after a Routing Engine switchover or restart.

As a workaround, you can configure static MAC addresses for aggregated Ethernet interfaces. Any external dependency, such as filtering of the MAC addresses that are assigned before the reboot, becomes invalid if the MAC address changes.

MPLS

- **Removal of SRLG details from the SRLG table only on the next reoptimization of the LSP**--If an SRLG is associated with a link used by an ingress LSP in the router, then on deleting the SRLG configuration from that router, the SRLG gets removed from the SRLG table only on the next reoptimization of the LSP. Until then, the output displays Unknown-XXX instead of the SRLG name and a nonzero srlg-cost of that SRLG for **run show mpls srlg** command.
- With the introduction of the multi-instance support for LDP-over-RSVP LSPs, you cannot enable MPLS on an interface that is already assigned to another routing instance. Starting in Junos OS Release 15.1R1, adding an interface that is part of another routing instance at the **[edit protocols mpls]** hierarchy level, throws a configuration error at the time of commit.

Network Management and Monitoring

- **Specified MIBS are not supported in Junos OS (MX Series)**--As of Junos OS Release 15.1, the following MIBS are not supported in Junos OS: CfmMepErrorCcmLastFailure, CfmMepXconCcmLastFailure, CfmMepCcmSequenceErrors, and ieee8021CfmMaCompNumberOfVids.

Subscriber Management and Services (MX Series)

- Junos OS Release 15.1R3 provides feature parity with the Junos OS Release 13.3R1 subscriber management feature set with the following exceptions:
 - Subscriber management is supported in the default logical system only.
 - TCP connections terminated by the router are supported for statically configured logical interfaces only. These connections are not supported on dynamically configured logical interfaces (for example, those using broadband edge dynamic-profiles).

- Bandwidth provisioning based on DHCP **option82** is not supported.
- The **dscp-code-point** statement at the **[edit class-of-service host-outbound-traffic]** hierarchy level is not supported.
- You cannot use the subscriber management configuration used for multicast **dynamic cos-adjust** with previous versions of Junos OS.
- Reverse outgoing interface (OIF) mapping, which enables the router to propagate the multicast state of the shared interface to the customer interfaces and enables per-customer accounting and QoS adjustments, is not available.
- Fast update filters for dynamic profiles, which you can use to incrementally add, remove, or update filter terms, are not supported.
- Lawful intercept of multicast traffic is not supported.
- Advisory speed reporting to a RADIUS server and to an L2TP network server (LNS) is not supported.
- Access Node Control Protocol (ANCP) is not supported.
- The use of the **per-unit-scheduler** statement at the **[edit interfaces interface-name]** hierarchy level is not supported for subscribers. Subscribers associated with this option do not appear.
- C-VLAN logical interfaces do not inherit the Ethernet Operation, Administration, and Maintenance (OAM) statuses from the associated S-VLAN logical interfaces.
- MS-DPCs are not supported.
- Static PPPoE interfaces are not supported.
- N-Way active targeting over aggregated Ethernet member links is not supported.
- Targeted distribution of static or dynamic interface sets over aggregated Ethernet is not supported.
- The **show ppp interface *interface-name* extensive** and **show interfaces pp0** commands display different values for the LCP state of a tunneled subscriber on the LAC. The **show ppp interface *interface-name* extensive** command displays STOPPED whereas the **show interfaces pp0** command displays OPENED (which reflects the LCP state before tunneling). As a workaround, use the **show ppp interface *interface-name* extensive** command to determine the correct LCP state for the subscriber.
- On MX Series routers, when you configure the **subscriber-awareness** statement on a service set by committing the **set services service-set *service-set-name* service-set-options subscriber-awareness** statement, the service sessions fail to create. To avoid this issue, on MX Series routers that support the Service Control Gateway solution, ensure that the Junos OS Mobility package software is installed on the router.

The Service Control Gateway solution is supported only in 14.1X55 releases. For Junos OS Releases 14.2, 15.1, and 16.1 ensure that the **subscriber-awareness** statement is not set.
- **Support for multicast group membership in Enhanced Subscriber Manager (MX Series)**— In Junos OS Release 15.1R3, enhanced subscriber management does not support the use of dynamic profiles for the static configuration of multicast group

membership for subscribers. Instead, subscribers must send an IGMP JOIN message to receive the multicast stream. More specifically, the following command is not supported in this release:

```
set dynamic-profiles client profile protocols igmp interface $junos-interface-name static
group 224.117.71.1
```

- **Dynamic provisioning in Layer 2 wholesaling (MX Series)**—Starting with Release 15.1R3, Junos OS does not support dynamic VLAN mapping into VPLS instances. (You can still configure static VLAN interface mapping to VPLS instances.) By extension, dynamic provisioning for Layer 2 wholesaling is also not supported in this release.

The following example shows the statements that are not currently available (**encapsulation vlan-vpls** and **family vpls** at the **[edit dynamic interfaces]** hierarchy level):

```
interfaces {
    "$junos-interface-ifd-name" {
        unit "$junos-interface-unit" {
            encapsulation vlan-vpls
                vlan-id "$junos-vlan-id";
            family vpls;
        }
    }
}
```

- **Preventing Link Aggregation Control Protocol link reversion in a scaled configuration (MX Series)**—By default, LACP link protection is revertive. This means that after the current link becomes active, the router switches to a higher-priority link if one becomes operational or is added to the aggregated Ethernet bundle. In a highly scaled configuration over aggregated Ethernet, we recommend that you prevent the router from performing such a switch by including the **non-revertive** statement at the **[edit chassis aggregated-devices ethernet lacp link-protection]** hierarchy level. Failure to do so may result in some traffic loss if a MIC on which a member interface is located reboots. Using the **non-revertive** statement for this purpose is not effective if both the primary and secondary interfaces are on the MIC that reboots.
- **Support for stacked IFL configurations (MX Series)**—Junos OS release 15.1R4 does not provide complete support for the interface variable, *\$junos-interface-ifd-name*, with stacked IFL configurations such as PPPoE. Juniper recommends using Junos OS release 15.1R5 or later if you need to reference more than one IFD within a dynamic profile (in other words, to stack IFLs in support of certain CoS configurations in conjunction with subscriber management).

System Logging

- **Text string deprecated in syslog messages that are converted to SNMP traps (M Series, MX Series, and T Series)**—In the syslog messages that are converted to SNMP traps for event policies, the "trap sent successfully" text string is deprecated.
- On MX Series routers, when you configure a rate limit for system log messages by setting the **message-rate-limit** statement for a multiservices interface, ensure that the **syslog host** option for that interface is configured. This configuration ensures that the system log statistics reflect the rate limit set for the interface.

VPNs

- **Default export EVPN policy has been removed (MX Series)**—Starting in Junos OS Release 15.1R5 and forward, the hidden default EVPN export policy statement (**evpn-pplb**) has been removed. To enable and configure load balance per packet for EVPN and PBB-EVPN, use the existing policy statements:
 - **set routing-options forwarding-table export evpn-pplb**
 - **set policy-options policy-statement evpn-pplb from protocol evpn**
 - **set policy-options policy-statement evpn-pplb then load-balance per-packet**



NOTE: To support EVPN multihoming, you must configure the **load-balance per-packet** statement.

- See Also**
- [New and Changed Features on page 80](#)
 - [Changes in Behavior and Syntax on page 131](#)
 - [Known Issues on page 170](#)
 - [Resolved Issues on page 192](#)
 - [Documentation Updates on page 312](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 320](#)
 - [Product Compatibility on page 329](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R6 for the M Series, MX Series, and T Series.

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

- [Forwarding and Sampling on page 171](#)
- [General Routing on page 172](#)
- [Infrastructure on page 175](#)
- [Interfaces and Chassis on page 175](#)
- [J-Web on page 175](#)
- [Layer 2 Ethernet Services on page 176](#)
- [MPLS on page 176](#)
- [Network Management and Monitoring on page 177](#)
- [Platform and Infrastructure on page 177](#)
- [Routing Protocols on page 180](#)

- [Services Applications on page 181](#)
- [Subscriber Access Management on page 181](#)
- [User Interface and Configuration on page 181](#)
- [VPNs on page 182](#)
- [Infrastructure on page 182](#)
- [Interfaces and Chassis on page 183](#)
- [Layer 2 Ethernet Services on page 184](#)
- [Multiprotocol Label Switching \(MPLS\) on page 184](#)
- [Network Management and Monitoring on page 185](#)
- [Platform and Infrastructure on page 186](#)
- [Routing Protocols on page 188](#)
- [Services Applications on page 190](#)
- [Subscriber Access Management on page 191](#)
- [User Interface and Configuration on page 191](#)
- [VPNs on page 191](#)

Forwarding and Sampling

- When VRRP is configured on MX Series routers with MPC/MIC-based MX Series interfaces, static MAC entries are installed on the Packet Forwarding Engine in the MAC database as part of MAC filter installation. The MIB walk on some object identifiers (OIDs) will trigger a walk over the MAC MIB entry (walk over the static MAC entries with no OIDs), resulting in an error message. During the walk, it is expected that no entries are read from static MAC database entries; however, the EODB is not set to indicate the MAC database walk has ended. This error log does not have any functional impact on the MIB walk: `mib2d[xxx]: MIB2D_RTSLIB_READ_FAILURE: check_rtsock_rc: failed in reading mac_db: 0 (Invalid argument) mib2d[xxx]: SNMP_GET_ERROR1: macStatsEntry getNext failed for interface: index1 ge-*/*/ (Invalid argument)` The following oid might trigger the issue: 1/ Rpf related oid 2/ AtmCos related oid 3/ Mac related oid , such as `jnxMacStatsEntry` 4/ PMon related oid 5/ `jnxSonetAlarmTable` 6/ `Scu` related oid 7/ `jnxCmRescueChg` 8/ `jnxCmCfgChgEventLog` 9/ `jnxIpv4AdEntReasmMaxSize`. [PR1042610](#)
- If a bandwidth-percent based policer is applied on an aggregated Ethernet (AE) bundle without the **shared-bandwidth-policer** configuration statement, traffic will hit the policer even if the traffic is not exceeding the configured bandwidth. As a workaround, configure the **shared-bandwidth-policer** configuration statement under the policer. [PR1125071](#)
- The "default-arp-policer" is applied to every relevant logical interface to rate-limit the Address Resolution Protocol (ARP) traffic. You can disable the "default-arp-policer" by running the command **set firewall disable-arp-policer**. Note that improper application results in the Routing Engine being overloaded with a bulk of ARP traffic, which leads to a typical Denial of Service (DoS) scenario. In this scenario, even after the "default-arp-policer" was disabled, it still affected the logical interface, such as after DUT reboot or when a new logical interface was created. [PR1198107](#)

General Routing

- Syslog **rate limit** value is always shown as zero in **show services service-sets statistics syslog**. This is a display issue. [PR900301](#)
- There is a 50 Kpps drop in performance related to new functionality. [PR935393](#)
- When the bcm0 interface goes down, the Routing Engine should switch over on the M320. [PR949517](#)
- If the ICMP echo response is sent with an incorrect sequence number, flow lookup passes and the counter gets incremented, but the packet is discarded by the ICMP ALG. [PR971871](#)
- When a MAC moves from one VTEP to another VTEP, it is not learned by the new VTEP until the old VTEP ages out this MAC. Traffic for this MAC is silently dropped or discarded until it ages out on the old VTEP. [PR988270](#)
- An NSX controller occasionally overrides an existing local MAC with a remote MAC of the same address. If a hardware VTEP in a Junos OS network detects such a condition (that is, it receives a remote MAC from the NSX controller that conflicts (matches) with an existing local MAC), the hardware VTEP in a Junos OS network accepts the remote MAC and stops publishing the local MAC to the NSX controller. [PR991553](#)
- This issue is seen when the configured global-mac limit is less than the interface mac limit and the same interface is configured with packet action. When the traffic is sent with higher packet rate, all the mac entries are learnt by the Packet Forwarding Engine. Routing Engine later trims this to the configured global-mac limit. When the traffic is sent with lower packet rate, the Routing Engine learns some what more than the configured global-mac limit and subjects the remaining packets (with newer macs) to the configured drop-action. [PR1002774](#)
- There is an existing optimization in RE kernel where the add IPCs of interface objects (IFD/IFL/IFF/IFA) are not sent to the FPCs (i.e. these IPCs get suppressed) when the corresponding IFD no longer has IFDF_PRESENT flag set. The idea is that since Chassisd has already removed this flag from the IFD, all daemons will start cleaning up the whole hierarchy and soon DCD will delete IFAs/IFFs/IFLs under it, before deleting the IFD itself. Kernel keeps track of which object's add IPC was suppressed for which FPC peer (it is a per object bit vector), and it suppresses the delete IPC too if the add was suppressed. This logic doesn't exist for RT and NH objects so sometimes it may happen that FPCs receives a NH IPC for which the parent IFL got suppressed in the kernel, hence it complains. It's a day-1 issue. There is no work around for this issue. These error messages are harmless as DCD would have deleted everything once scheduled. [PR1015941](#)
- In an MPLS L3VPN scenario, if the ingress is MX Series with MPCs/MICs-based line card and the egress is hosted on M120, M7i/M10i with E-CFEB/Enhanced, E3-FPC in M320 or MX ADPC/E line cards, the packets will be truncated if an MPLS experimental (EXP) rewrite rule is applied with "mpls-inet-both-non-vpn" or "mpls-inet-both" configuration on the egress and the **chained-composite-next-hop** configuration statement is configured. [PR1018851](#)

- Time taken to reboot T series boxes has gone up. T Series(Standalone) 14.2 - 3 minutes 39 seconds 15.1 - 4 minutes 18 seconds (Difference - 40 seconds) TX Matrix (Multichassis) 14.2 - 5 minutes 17 seconds 15.1 - 7 minutes 18 seconds (Difference - 2 minutes) [PR1049869](#)
- On MX Series routers where FPCs or MPC or MIC line cards(MX240, MX480, and MX960) are supported that run Enhanced Layer 2 Software (ELS), when an interface is removed from a private VLAN (PVLAN) and then added back. The corresponding MAC entry might not be deleted from the Ethernet table. [PR1036265](#)
- When using the 'mpls-ipv4-template' sampling template for non-IP traffic encapsulated in MPLS, log messages such as this one can be seen frequently (depending upon the rate of traffic, which could range from a few messages up to 3000 messages per minute): **Router-re0 : %DAEMON-3: (FPC Slot 2, PIC Slot 0) ms20 mspmand[171]: jflow_process_session_close: Could not get session extension: 0x939d53448 sc_pid: 5.** Depending upon the frequency of the messages per second, the eventd (daemon) utilization can shoot up processing of these syslogs at the Routing Engine. Eventually high CPU utilization is observed at the Routing Engine, which can be checked by the commands **show chassis routing-engine** or the **freebsd "Top"** command under the shell. [PR1065788](#)
- The customer configuration is logically forming a loop. BGP tries to use inet.3 to perform a next-hop resolution. With the current configuration (next-table), BGP is asking to perform resolution from inet.0 again. From the standpoint of the packet flow, the packet lookup happens on inet.0, and then gets a "next-table inet.0" instruction. This means starting from inet.0 to do the lookup again, which is Step 1 of the lookup. This is causing the loop. [PR1068208](#)
- With MPC5E, high temperatures cause the fan to spin continuously. [PR1070346](#)
- On MX Series platform with MS-MPC/MS-MIC, memory leaks will be seen with `jnx_msp_jbuf_small_oc` object, upon sending millions of PPTP control connections (3-5M) alone at higher connections per second (> 150,000 CPS). This issue is not seen in up to 50,000 control connections at 10,000-30,000 connections per second. [PR1087561](#)
- In Service profile, we have same variable used by filter and CoS, for example, in profile RLInternet, variable OutFilter is used by out filter and CoS TCP. This is wrong concept as filter and CoS should have different variables. To fix the issue, we need to do the following: 1. In service profiles, add variable for CoS TCP, for example, in profile RLInternet, add OutTcp for TCP 2. In JSRC, add value for TCP variable OutTcp, for example, `?OutTcp:any="1M"` 3. In JSRC, change value of out filter to distinguish from TCP, for example, `'OutFilter:any="1M-out"` 4. Modify out filter name, for example, from "1M" to "1M-out" [PR1154982](#)
- The speed configuration statement `auto-10m-100m` allows you to autonegotiate the maximum speed to 100 Mbps. [PR1155196](#)
- The stacked logical interface and the underlying logical interface cannot be part of the same iflset. [PR1162805](#)
- On T Series multichassis platform, when offline and then online the LCC from SCC (for example, executing the CLI command **set chassis lcc 0 offline**, and then executing

delete set chassis lcc 0 offline) in quick successions (that is, within the timeout setting for peer to reconnect, 60 seconds, which is not configurable), kernel replication error "ENOENT" might occur, which can cause ksyncd to crash and in thus trigger a live vmcore. Additionally, this is a timing issue and LCC offline followed by online within 60 seconds is the only known trigger so far. As a workaround, on the safer side, it is recommended to online the LCC after 120 seconds. [PR 1108048](#)

- When two or more Ethernet VPN (EVPN) PEs are connected to a multi-homed CE, after switching the router-id between the PEs, the non-designated forwarder (DF) PE will forward Broadcast, unicast unknown, and multicast (BUM) traffic back to the Ethernet Segment Identifier (ESI). [PR1108107](#)
- The rdd (a daemon used by MS-DPC/MS-MPC/AE) might crash after performing certain operations on a dual Routing Engine MX Series platform with AE interface configured and a non-Trio FPC installed. [PR1188832](#)
- In case of local source and with asm MoFRR enabled, the default MDT traffic loops back to the originating router on the MoFRR backup interface, thereby causing continuous IIF_mismatches. With the current MoFRR code, since the source is local, the SPT BIT is set by default; hence an (S,G,rpt) PRUNE is sent out of the MoFRR active interface. But an (S,G,rpt) PRUNE is not sent out of the MoFRR Backup interface (missing code). An (S,G,rpt) PRUNE should be sent over the MoFRR backup path also (if there is already an (S,G,rpt Prune) going out of the MoFRR active path) in order to avoid IIF_Mismatches. [PR1206121](#)
- On very rare occasions, offlining a MIC-3D-16CHE1-T1-CE MIC generates FPC core file. There is no workaround for this except to upgrade to an image with this fix present. [PR1223277](#)
- This issue is seen in Junos OS Release 15.1R1 and Junos OS Release 16.1X70-D10. After setting the following command and rebooting the router, you see some error logs on terminal display and messages log. This error message looks like a cosmetic issue: **=== command === set system syslog user * any critical commit** . After rebooting the router, this message is seen **=== error message === router-re0 /kernel: GENCFG: op 34 (CLKSYNC blob) failed; err 7 (Doesn't Exist) <<<<<<< [PR1223518](#)**
- An incorrect PE device is being attached to an ESI when the router receives two copies of the same AD/ESI route (for example, one through eBGP and another one received from an iBGP neighbor). This will causes partial traffic blackhole and stale MAC entries. You can confirm the issue by checking the members of the ESI: **user@router> show evpn instance extensive ... Number of Ethernet segments: 5 ESI: 00:13:78:00:00:00:00:00:01 Status: Resolved Number of remote PEs connected: 3 Remote PE MAC label Aliasing label Mode 87.233.39.102 0 0 all-active 87.233.39.1 200 0 all-active <<<< this PE is not part of the ESI 87.233.39.101 200 0 all-active.** [PR1231402](#)
- DNS server IP addresses are not present in the output of **show subscribers extensive** for DHCP subscribers in case DNS configuration is provided from the access-profile or pool. If such data is provided from RADIUS, the output is correct. The issue is cosmetic: DNS addresses are provided to subscribers. [PR1237525](#)
- In subscriber management environment, starting from Junos OS Release 15.1R4, new generation architecture is used (aka Tomcat). In these new releases, if dynamic VLAN profile does not have IFF configuration (e.g. either family pppoe or family inet), even

if firewall filter configuration is present in the dynamic profile it won't be programmed properly at the forwarding complex (PFE). [PR1264367](#)

- When RSTP protocols configuration is disabled and enabled on the distribution switch using CLI, some times it is found that the ARP request packets are not egressing out on aggregate L2 interface where host is connected for the inter-VLAN traffic ingressing on other aggregate L2 interface from an access switch. The L3 interface (irb) is configured for VRRP on the distribution switch under this topology. [PR1265471](#)
- It is possible to see a bbe-smgd core under certain boundary conditions on the standby Routing Engine with certain specific configurations. Since the core is on the standby no disruption in service is expected and system recovers from this condition. [PR1267646](#)

Infrastructure

- On M10i or M7i router, the Routing Engine goes to # Prompt or db> prompt after setting mirror-flash-on-disk. [PR1260268](#)

Interfaces and Chassis

- On dual Routing Engine platforms, when adding the logical interfaces (IFLs) and committing, due to the device control process (dcd) on the backup Routing Engine, it might fail to process the configuration and keep it in the memory. In some cases (not happening all the time), it might be observed that the memory of the dcd keeps increasing on the backup Routing Engine. [PR1014098](#)
- After changing the MTU on the physical interface, on the static VLAN demultiplexing interface above the physical interface, an IPv6 link local address is not assigned. [PR1063404](#)
- During failure notification state machine, CFM does not correctly transit from DEFECT CLEARING state to RESET once the error indication has been cleared. As a consequence all the forthcoming errors will be considered post errors and will be reported right away without incurring the fngAlarmTime. This is a cosmetic problem. [PR1096346](#)
- The neighbor-ship is not created on IRB interface. [PR1198482](#)
- T3 interfaces configured with "compatibility-mode digital-link" might fail to come up due incorrect subrate. To verify, **show interfaces t3-0/0/0 extensive** will display the subrate. **DSU configuration: Compatibility mode: Digital Link, Scrambling: Enabled, Subrate: 4195338 Kbps <<<< // expected result: DSU configuration: Compatibility mode: Digital Link, Scrambling: Enabled, Subrate: Disabled** [PR1238395](#)

J-Web

- When you open a J-Web interface session using HTTPS, enter a username and a password, and then click the Login button, the J-Web interface takes 20 seconds longer to launch and load the Dashboard page than it does if you use HTTP. [PR549934](#)
- When the J-Web interface is launched using HTTPS, the time shown in the View Events page (Monitor >Events And Alarms > View Events) differs from the actual time in the switch. As a workaround, set the correct time in the box after the J-Web interface is launched. [PR558556](#)

Layer 2 Ethernet Services

- IPv4 and IPv6 long Virtual Router Redundancy Protocol (VRRP) convergence delay and unexpected packet loss might happen when MAC move for the IRB interface occurs (e.g. when flapping the Layer 2 interface which is the under-interface of IRB on master VRRP). [PR1116757](#)
- This issue occurs when you are running LACP between Juniper Networks and Cisco devices with different timers (Juniper fast and Cisco slow) on both sides. On the Cisco device side, it takes almost 90 seconds to bring the interface down from the bundle. When one interface is removed from the LAG on the Juniper Networks device side, the lead to the Cisco device needs to time out to bring the interface down from the bundle. This results in unexpected behavior outage on the network. [PR1169358](#)
- SNMP trap implemented in case of a single input feed fails for MX960 routers equipped with high-capacity PEMs. [PR1189641](#)
- When MSTP is configured under routing-instance, both the primary and standby VPLS pseudowires are stuck in ST state due to a bug in the software. [PR1206106](#)
- After changing the underlying physical interface for a static VLAN demultiplexing interface, the NAS-Port-ID is formed still based on the previous physical interface. [PR1255377](#)

MPLS

- If the **edit-protocols-mpls-traffic-engineering** configuration statement is configured, you cannot downgrade from Junos OS Release 14.2 to an earlier release. In order to downgrade, you must delete the **traffic-engineering stanza** and reconfigure it after downgrade. [PR961717](#)
- When using **mpls traffic-engineering bgp-igp-both-ribs** with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit ABR when there are two or more such ABRs. This causes interdomain RSVP LSPs to break. RSVP LSPs within same area are not affected. As a workaround, you can either run RSVP only on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/ISIS L2 routers, or avoid LDP completely and only use RSVP. [PR1048560](#)
- The multi-instance RSVP feature might not work in specific scenarios on MX Series devices with MPC cards when the core facing interface of the virtual routing instance and VPLS pseudowire termination is simultaneously configured on the master instance. As a workaround, configure the `?import-label-route?` statement at the `[edit routing-instances <routing instances name >protocols vpls]` hierarchy level to resolve this issue. [PR1080714](#)
- In BGP prefix-independent convergence (PIC) edge scenario, when the ingress route (the primary route) fails, due to the fact that LDP may fail to send the session down event to PFE correctly, the PFE may still use the primary path to forward traffic until (in some cases, 3- 5 seconds for 30k prefixes) the global convergence is completed by the interior gateway protocol (IGP). In addition, the issue may also be seen when the "delay-delete" knob is configured, in this scenario, the session down event may

get sent to the PFE correctly, however, due to local reversion, the primary path may also be chosen as forwarding path when it is deleted. [PR1097642](#)

- Benign error messages are generated; they can be ignored. [PR1136033](#)
- In some Inter-op scenario with cisco, sometime a new label advertised with withdraw the old label by cisco. Under such scenario, Junos OS will reject the new label advertised (as per RFC3036 behavior) Below mentioned logs will be generated in such event::
Line 408105: Mar 14 14:00:21.716559 LDP: LabelMap FEC L2CKT NoCtrlWord ETHERNET VC 40347 label 53 - received unsolicited additional label for FEC, releasing new label
This is an expected behavior as per RFC 3036. Junos will implement (17.1 onwards) behavior defined under RFC5036 in future releases as per which new label will be accepted and old label will be discarded. [PR1168184](#)
- In certain scenarios, the entropy label value being generated might not provide a good load-sharing result. [PR1235258](#)
- Routing Protocol process (RPD) might stop running unexpectedly if a static MPLS LSP is moved from one routing instance to another routing instance in one single configuration change with one single commit. RPD will need a manual restart with "restart routing". [PR1238698](#)

Network Management and Monitoring

- Eventd uses event library for signal handling. Although rare, a core file can occur due to a race condition synchronization issue in event library while handling signals. Event library is not signal safe and thus it is vulnerable to such issues. Eventd handles different kinds of signals (Through Signal Handlers) - SIGHUP (On Commit), - SIGTERM (On killing eventd) - SIGCHLD (on termination of event script execution) - SIGUSR1 & SIGUSR2 (on log rotation) If one signal handler is preempted by another signal handler, then it can adversely affect Wait List structures (and this generates core files). This can happen when eventd receives a new signal, while it is processing another signal. [PR1122877](#)

Platform and Infrastructure

- When next-ip is defined as the action and there is no ARP for the IP address specified under next-ip, the traffic is not forwarded. A manual ping needs to be initiated for things to work. [PR864861](#)
- When there is huge logical interface (IFL) scaling on the aggregated Ethernet interface (500 or more) with more than 32 member links and when all FPCs are restarted one by one, followed by member link addition to the link aggregation group (LAG), the state dependency evaluation in the kernel will take a long time. Because of the scale involved, the FPCs will not get all the states from the Routing Engine. Because it is a pretty uncommon sequence of events, the likelihood of this happening is rare. [PR938592](#)
- The overhead values need to be represented with 8 bits to cover the range "-120..124", but the microcode is only using the last 7 bits. [PR1020446](#)
- When TCP authentication is enabled on a TCP session, the TCP session might not use the selective acknowledgement (SACK) TCP extensions. [PR1024798](#)

- IPv6 packet loss and traffic degrades occur because MX Series routers have a restrictive rate limit on ICMPv6 packets that are too big. [PR1042699](#)
- Once the Traffic Offload Engine thread is stalled due to a memory error at the lookup chip, all statistics collected from the interfaces hosted by this Packet Forwarding Engine are not updated anymore. [PR1051076](#)
- In configurations with IRB interfaces, during times of interface deletion, such as an FPC reboot, the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system re-converging on the expected state. [PR1054798](#)
- Parity errors might be seen in pre-classifier engines. Such errors are not reported. As such, these packets will be silently discarded. Common symptoms are an increase in value of the "Input DA rejects" counter on the affected interface or misclassification of input packets. [PR1059137](#)
- Juniper VSA length above 2K bytes is not supported. Using authorization parameters above this length would result in wrong authorization setting for the user. [PR1072356](#)
- When deleting some uncommitted configuration on active RE, the rpd process on backup RE might restart due to "Unable to proceed with commit processing due to SIGHUP not received. Restarting to recover". [PR1075089](#)
- In XM-based multi-LU systems (MX Series platform with MPC3E/MPC4E/MPC5E/MPC6E/NG-MPC3/NG-MPC2 or T4000 with T4000-FPC5-3D linecard), when you have multiple LUs representing the same Packet Forwarding Engine complex and designate the BFD processing to a dedicated LU (LU 0), called as anchor LU, the rest of the LUs (LU 1, LU 2, LU 3) are called non-anchor LUs. When the Inline BFD packets punt from the non-anchor LU to the anchor LU, interface-group is not populated in the packet context, so the packets might not be matched by the related filter term. [PR1084586](#)
- On MX Series with MPCs/MICs-based platform, the Bidirectional Forwarding Detection (BFD) session over Integrated routing and bridging (IRB) interface, with a static client, might not come up with a Virtual Router Redundancy Protocol (VRRP) configuration. [PR1085599](#)
- Under large-scale setup, VPLS MAC might not be aged out from the remote Packet Forwarding Engine when the local Packet Forwarding Engine is MPC3/MPC4/MPC3E/MPC4E. Unknown unicast frames flood will be seen on local Packet Forwarding Engine. [PR1099253](#)
- Service chaining of inline software with NAT is not supported. Currently, when you commit configuration with software rule and NAT rules under same service-set on SI interface, commit goes fine. [PR1136717](#)
- We support maximum of 1024 Software concentrators with Inline-6rd. When we configure more than 1024 software concentrators (for example, 1025), you see a commit error message that is not very informative. [PR1153092](#)
- On MX Series with MPCs/MICs platform, the system might try to access a NULL pointer returned from hardware state lookup of the logical interface when the system runs out of memory. This can result in an FPC crash. [PR1163606](#)

- The delegated BFD session over AE interface failed to come up after FEB switchover with FEB redundancy group (1:1 and 1:N). [PR1169018](#)
- Internal fabric header corruption on MX Series Packet Forwarding Engines can lead to packet corruption on the egress Packet Forwarding Engine. This PR effort is to protect the fabric header coming to the egress MX Series Packet Forwarding Engine with a fabric CRC check. This approach helps to avoid wedges caused due to corrupted fabric headers. This PR adds a 32-byte CRC to each fabric packet sent from the ingress Packet Forwarding Engine to the egress Packet Forwarding Engine. On the egress Packet Forwarding Engine this CRC hash is validated; if the check fails the packet will be dropped. Because corrupted packets are dropped by this method, it can avoid potential ASIC wedges caused by bad hardware sending corrupted packets in the chassis. A new CLI command is added under chassis: `set chassis fabric-header-crc-enable`. Once this configuration statement is configured and committed, it will display an error message to reboot the box. It's recommended to configure this configuration statement in a service window with all traffic drained from the box. An immediate reboot is recommended once the configuration statement is configured. Inability to reboot the box can lead to unexpected packet drop behavior including wedges. All FPC types (except Ichip DPCs) on MX240/480/960/2010/2020 are affected by this bug. MX80/MX104 are not impacted by this bug, because they have a single Packet Forwarding Engine and this command is not supported on those platforms. Type-5 FPCs on T4000 and TXP+ are also impacted and this configuration statement will be supported on these platforms. [PR1170527](#)
- With MAC accounting feature (configuration `ethernet-switch-profile mac-learn-enable`) configured on an interface of MX Series with MPCs/MICs-based FPC, the limit of the MAC database might be reached and the FPC might crash. [PR1173530](#)
- Several files are copied between Routing Engines during 'ffp synchronize' phase of the commit (for example `/var/etc/mobile_aaa_ne.id`, `/var/etc/mobile_aaa_radius.id`). These files are copied even if there was no corresponding change in the configuration thus unnecessarily increasing commit time. [PR1210986](#)
- IPv6 traffic that is learned on an L2/bridge/multilink interface and has been traversed through MPLS core random packets might get classified incorrectly by the fabric which leads to packet loss. [PR1223566](#)
- In affected release with bridge over GRE feature configured, the traffic failed to send to the GRE underlying interface once bouncing the gr-interface. [PR1255706](#)
- This issue is specific to a router running Junos OS Releases 15.1 or 15.1F, which also have VRRP with PIM configured. When a router's interface VRRP mastership switches from VRRP master to backup, the router continues to use VRRP virtual MAC address (VMAC) for the source address of all Ethernet packets leaving its interface. [PR1257477](#)

Routing Protocols

- The multicast next hop **show multicast nexthop** shown for master and backup Routing Engine for the same flow could be different if the next hop is hierarchy MCNH. When doing NSR switch, however, there is no traffic loss caused by this show difference. [PR847586](#)
- In rare cases, rpd might write a core file with the signature **rt_notbest_sanity: Path selection failure on** The core is "soft", which means there should be no impact to traffic or routing protocols. [PR946415](#)
- For FEC 129 VPLS (also known as LDP VPLS with BGP-based autodiscovery), if abandoned VRF and VPLS instances are left after all of the other pieces of configuration are removed, and the BGP protocol is deactivated in the master instance, the rpd process might crash continuously when you commit a new configuration. As a workaround, you should remove all the unused VRF and VPLS instances. [PR1006689](#)
- Scaled configurations toggling from 64-bit to 32-bit rpd at the same time that Rosen MVPN routing instances are deleted can result in a kernel core file on the backup Routing Engine. [PR1022847](#)
- The static/static access routes pointing to an unnumbered interface are getting added in the routing table even if the interface is down. In this case, if graceful Routing Engine switchover (GRES) is disabled, these type of routes will never be added in the routing table after Routing Engine switchover. [PR1064331](#)
- When multiple addresses are configured on an interface, if the interface has "interface-type p2p" configured under OSPF and the router does not receive any OSPF packets from one of the IFAs, the OSPF state will not go down for the corresponding adjacency. It should have no impact on route learning, but it might cause confusion for troubleshooting, when peering with Cisco devices, which have multiple addresses configured as secondary addresses. [PR1119685](#)
- A few seconds of traffic loss is seen on some of the flows when the PE-CE interface comes up and the PE router starts learning 70,000 IPv4 prefixes and 400 IPv6 prefixes from the CE router during L3VPN convergence. [PR1130154](#)
- In a multicast environment, when the rendezvous point (RP) is a first-hop router (FHR) and it has MSDP peers, when the rpf interface on the RP changes to an MSDP-facing interface, because the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)
- When applying add-path prefix-policy to neighbor level, all neighbors are separated into different update groups. This is not expected behavior. There is no service impact. But if all the neighbors are configured under one peer group, if there is a huge number of peer groups, the scaling/performance will go down. [PR1137501](#)
- Generate route does not inherit the next-hop from the contributing route in L3VPN when the contributing route is learned through MP-BGP. The next hop remains as rejected for the generated route. [PR1149970](#)
- Junos OS marks hidden routes with a negative route preference. The router policy-statement explicitly sets a preference value for the BGP routes, including the hidden routes, while the routes are imported into the routing instance table. The negative

route gets overridden, but AS_PATH loops are not checked within the VRF context. This results in the hidden routes becoming active in the routing instance table. [PR1165781](#)

- See the following topology. If the opposite router's interface "A" is down because of a disable, deactivate, or delete configuration or a transmission issue, BFD timeout detection might be greatly delayed. Topology +-----+ | DUT | OSPF | |-----+ +-----+ | A | | | | | +-----+ OSPF(p2p) | | R2 | bfd | | | | +-----+ | | V intf A | | +-----+ | | R1 |-----+ | | OSPF +-----+ [PR1183353](#)
- Here are the results when L1 is disabled for Lo0: **{master}[edit] user@router# run show isis interface IS-IS interface database: Interface L CirID Level 1 DR Level 2 DR L1/L2 Metric lo0.0 3 0x1 Disabled Passive 0/0**. Here are the results when L2 is disabled for Lo0: **{master} user@router> show isis interface IS-IS interface database: Interface L CirID Level 1 DR Level 2 DR L1/L2 Metric lo0.0 3 0x1 Passive Disabled 0/0**. [PR1202216](#)
- In the context of a large number of configured VPNs, routes changing in the midst of a bgp path-selection configuration change can sometimes lead to an rpd core file. This core file has been seen with the removal of the **always-compare-med** option. [PR1213131](#)

Services Applications

- In the NAT environment, the jnxNatSrcPoolName OID is not implemented in jnxSrcNatStatsTable. [PR1039112](#)
- This issue occurs in a typical scaling scenario, where there are a high number of simultaneous updates from the Routing Engine to the Packet Forwarding Engines. If there is a request for packet mirroring, the message might get delayed because of other pending messages, resulting in an overall delay in mirroring of the requested flows. [PR1244849](#)
- This is a limitation in Junos OS 15.1 releases and later LI implementation. To enable this support it will require considerable effort to implement it. In current build, if customers really want to disable the drop policy, then they can delete the LI service via DTCP delete and then add the LI service back by DTCP add without drop policy. [PR1252079](#)

Subscriber Access Management

- Call rate performance might be impacted under heavy load if there are large numbers of small linked address pools because of a bug in the allocation traversal algorithm. [PR1264052](#)

User Interface and Configuration

- Selecting the Monitor port for any port in the Chassis Viewer page takes the user to the common Port Monitoring page instead of the corresponding Monitoring page of the selected port. [PR446890](#)
- On the J-Web interface, the Configure > Routing > OSPF > Add > Interface Tab is showing only the following three interfaces by default: pfh-0/0/0.16383, lo0.0, lo0.16385. To overcome this issue and to configure the desired interfaces to associated ospf area-range, perform the following operation in the CLI: **set protocols ospf area 10.1.2.5 area-range 12.25.0.0/16, set protocols ospf area 10.1.2.5 interface fe-0/3/1**. [PR814171](#)

- For configure > clitools > point and click > system > advanced > deletion of saved core, the No option is not available in J-Web. [PR888714](#)
- Basic value entry format error check is not present in Configure-->Security-->IPv6 Firewall Filters, but the same is present in IPv4 Firewall Filters. But it will throw error when try to commit the wrong format data entered. [PR1009173](#)

VPNs

- (Refer to release note of [PR 535844](#)). The default BGP extended community value used for MVPN IPv4 VRF Route Import (RT-Import) should be modified to meet the IANA-standardized value. That is, the behavior of the configuration **mvpn-iana-rt-import** should become the default and the **mvpn-iana-rt-import** configuration should be deprecated. [PR890084](#)
- On the backup Routing Engine, the routing process can sometimes crash while it is performing block label allocations for L2CKT/L2VPN applications. This can typically occur at a time of lot of churn, say, several routing-instances and deactivated and re-activated coupled with DUT and Peer router's restart. [PR119684](#)
- For a next-generation multicast VPN (NG-MVPN) using ingress replication provider tunnels, if both IPv4 and IPv6 are configured, when the receiver PE router advertises different labels for IPv4 and IPv6 in the type-1 BGP route, the source PE router will create two provider tunnels to carry both IPv4 and IPv6 traffic, causing duplicated multicast traffic. [PR1128376](#)
- The impact on customers will be less since this is only happening on doing vrf-localization. The routes will be deleted and re-added again when this is done and there is bound to be some loss. In the current scenario however, after the routes deleted, it was taking slightly more time to add it back and get stabilized. When the script checks for the route after vrf-localization is done, it was not available, but it comes up after sometime and everything should work fine. [PR1264366](#)

Infrastructure

- The syslog messages are observed when one of the following CLI commands is executed: **system syslog file messages kernel any** or **system syslogfile messages any any**. These syslog messages do not indicate any functionality, breakage, or impact. If you need to enable "anyany", then you would need to skip these logs with an appropriate match condition. [PR1239651](#)
- When the configuration statement **set system ports console log-out-on-disconnect** is enabled, the Junos OS eventd process (daemon) blocks the console-open(). However, during this stage with the syslog console configured (always logs on console), any logging continues even if the console session is ended. When the console logging continues to be in the waiting status, the eventd syslog rotation freezes and some processes directly involved in logging in to the system would also go into the wait status, causing undesirable behavior. [PR1253544](#)

- The kernel might fail to finish all input or output before shutdown during updating. And the upgrade might not succeed with the following reason: "Could not find installation package". [PR1298749](#)
- Juniper Networks Routing Engines with HAGIWARA CF card installed, after upgrade to Junos OS 15.1 and later releases, reports the following failure message (every 30 minutes): `root@router> show log messages| match ada1 Jan 15 21:18:28 router smartd[4357]: Device: /dev/ada1, failed to read SMART Attribute Data Jan 15 21:48:28 router smartd[4357]: Device: /dev/ada1, failed to read SMART Attribute Data Jan 15 22:18:28 router smartd[4357]: Device: /dev/ada1, failed to read SMART Attribute Data <...>` Issue related with Smartd daemon reading the SSD attributes with HAGIWARA SYS-COM MFD10 Compact Flash. Routing Engine 0 REV 05 740-031116 xxxxxxxxxxxx RE-S-1800x4 <...> ada1 3671 MB HAGIWARA SYS-COM MFD10 xxxxxxxxxxxx Compact Flash. [PR1333855](#)

Interfaces and Chassis

- MX Series chassis control interrupt storm might be falsely reported when a field replaceable unit (FRU) is removed, inserted, or FPM button is pushed. A FRU might not be recognized or booted, resulting in chassis operational failure. [PR823969](#)
- The internal management Ethernet interfaces (em-) might fail auto-negotiation after a reboot if one of the em interfaces is in a link down condition. It might hold back the linecards from booting after the Routing Engine reboots and kernel memory keep increasing. Any platform using "em" interfaces might experience a kernel memory leak of "devbuf" buffers when the interface is administratively marked up, but the physical link is down. Once the kernel "devbuf" allocations exceeds an internal threshold, warnings messages are recorded in the message file: **kernel: kmem type devbuf using <X> K, exceeding limit <Y> K**. These logs remains until the Routing Engine is restarted. See PR workaround for how to stop further kernel "devbuf" memory from leaking. [PR829521](#)
- Online insertion and removal (OIR) is not supported on a PIC (PD-4XGE-XFP) currently. When pulling out a PIC (PD-4XGE-XFP) from an FPC that is not offline, the FPC crashes generating a core file. [PR874266](#)
- On dual Routing Engine platforms, when adding the logical interfaces (IFLs) and committing, the device control process (dcd) on the backup Routing Engine might fail to process the configuration and keep it in the memory. In some cases, the memory of the dcd keeps increasing on the backup Routing Engine. [PR1014098](#)
- Configuring ODU FRR under otn-options for 2x100G DWDM PIC is an unsupported command on MX Series routers. Incorrectly adding such a configuration could result in an FPC crash and restart. [PR1038551](#)
- After changing the MTU on the physical interface, an IPv6 link local address is missing on the static VLAN demux interface. [PR1063404](#)
- During failure notification state machine, CFM does not correctly transit from DEFECT CLEARING state to RESET once the error indication has been cleared. As a consequence all the forthcoming errors will be considered post errors and are reported immediately without incurring the fngAlarmTime. This is a cosmetic problem. [PR1096346](#)

- DWDM 1x100G MIC FPC restart with next generation MPC core file is generated. An interruption is observed when the forward error correction(FEC) mode is changed resulting in an FPC crash. [PR1232392](#)
- In a VPLS multihoming scenario, the CFM packets are forwarded over the standby PE device link, resulting in duplicate packets or a loop between the active and standby link. [PR1253542](#)
- In subscriber management scenario with dynamic demultiplexing interfaces (DEMUX) configured, in the case where subscribers belonging to one aggregated Ethernet interface are migrated to a newly configured aggregated Ethernet interface, subscribers might fail to access the device after deleting the old aggregated Ethernet configuration. [PR1322678](#)

Layer 2 Ethernet Services

- IPv4 and IPv6 long VRRP convergence delay and unexpected packet loss might occur when a MAC move for the IRB interface occurs (for example, when flapping the Layer 2 interface that is the underlying interface of IRB on the master VRRP). [PR1116757](#)
- When MSTP is configured under a routing instance, both the primary and standby VPLS pseudowires get stuck in ST state due to a bug in the software. That has been fixed and now the PW status is set correctly. [PR1206106](#)
- On MX240, MX480, and MX960 platforms, due to a resources contention during multiple commit processes, the kernels might display I2C bus errors. [PR1298612](#)

Multiprotocol Label Switching (MPLS)

- When the **edit->protocols->mpls->traffic-engineering** statement is configured, the user cannot downgrade from a Junos OS Release 14.2 to prior releases. In order to downgrade, the user must delete the traffic-engineering stanza and re-configure it after downgrade. [PR961717](#)
- When using **mpls traffic-engineering bgp-igp-both-ribs** with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit area border router (ABR) when there are two or more such area border routers (ABRs). This causes interdomain RSVP LSPs to break. RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. [PR1048560](#)
- The multi-instance RSVP feature might not work in specific scenarios on MX Series devices with MPC cards when the core facing interface of the virtual routing instance and VPLS pseudowire termination is simultaneously configured on the master instance. As a workaround, configure the **import-label-route** statement at the **[edit routing-instances <routing-instance-name> protocols vpls]** hierarchy level to resolve this issue. [PR1080714](#)
- In a BGP prefix independent convergence (PIC) edge scenario, when the ingress route (the primary route) fails, because the LDP might fail to send the session down-event to the Packet Forwarding Engine correctly. The Packet Forwarding Engine might still use the primary path to forward traffic until (in some cases, 3-5 seconds for 30,000

prefixes) the global convergence is completed by the interior gateway protocol (IGP). This issue might also be seen when the **delay-delete** CLI command is configured. In this scenario, the session-down event is sent to the Packet Forwarding Engine. However, due to local reversion, the primary path might also be chosen as the forwarding path when it is deleted. [PR1097642](#)

- Benign error messages can occur when bringing up new interfaces or line cards; they can be ignored. [PR1136033](#)
- The routing protocol process (rpd) might stop running unexpectedly if a static MPLS LSP is moved from one routing instance to another routing instance in one single configuration change with one single commit. The rpd needs a manual restart with **restart routing**. [PR1238698](#)
- In CE-CE setup, traffic loss might be observed over secondary LSP on primary failover. [PR1240892](#)
- When MPLS builds the next hop for an mpls.0 route for the scenario with IDP over RSVP LSP over bypass tunnel and the IDP label is implicit-NULL, the label stack constructed for the next hop might be incorrect, with an invalid bottom label value of 1048575. [PR1270877](#)
- The **show mpls container-lsp** output does not show any egress LSP until the Enhanced FRR is enabled for these egress LSPs. [PR1314960](#)

Network Management and Monitoring

- Eventd uses event library for signal handling. The core file is generated because of the race-condition/synchronization issue in event library while handling signals. Event library is not signal safe and thus it is vulnerable to such issues. Eventd handles different kinds of signals (through signal handlers)
 - SIGHUP (on commit)
 - SIGTERM (on killing eventd)
 - SIGCHLD (on termination of event script execution)
 - SIGUSR1 & SIGUSR2 (on log rotation)

If one signal handler is preempted by another signal handler, then it might mess up waitList structures (and this can generate the core file). This might happen when eventd receives a new signal while processing another signal. This situation might not occur every-time and hence generating of core file is observed rarely. [PR1122877](#)

- Traps are sent as AgentX messages type (AGENTX_MSG_NOTIFY) from subagent to master agent. Subagent expects a response in the form of an acknowledgement from snmpd after sending these AGENTX_MSG_NOTIFY messages upstream. If an ACK is not received from snmpd within 1 second (current timeout value), the sub-agent resends the trap. After router reboots or GRES a lot of upstream communication is triggered from subagent to snmpd (traps/mib registration messages). This time the snmpd process might not be able to send the downstream ACK within 1 second period. This might trigger the subagent to resend trap and this is seen as a duplicate trap on the NMS. [PR1164848](#)

Platform and Infrastructure

- When next-ip is defined as the action and there is no ARP for the IP address specified under next-ip, the traffic is not forwarded. A manual ping needs to be initiated for things to work. [PR864861](#)
- When there is huge logical interface (IFL) scaling on aggregated Ethernet interfaces (500 or more) with more than 32 member links and when all FPCs are restarted one by one, followed by member link addition to the link aggregation group (LAG), the state dependency evaluation in the kernel takes a long time when the scale involved is given. As a result, the FPC's does not get all the states from the Routing Engine. This is an uncommon sequence of events or conditions. [PR938592](#)
- The overhead values need to be represented with 8 bits to cover the range "-120 to 124", but the microcode is only using the last 7 bits. [PR1020446](#)
- When TCP authentication is enabled on a TCP session, the TCP session might not use the selective acknowledgement (SACK) TCP extensions. [PR1024798](#)
- When using **show | compare** method to commit, part of configuration might be treated as noise and return syntax error. [PR1042512](#)
- Once the Traffic Offload Engine (TOE) thread is stalled due to a memory error at the lookup chip, all statistics collection from the interfaces hosted by the Packet Forwarding Engine are not updated anymore. [PR1051076](#)
- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. [PR1054798](#)
- On MX Series routers, parity memory errors might occur in pre-classifier engines within an MPC. Packets are silently discarded because such errors are not reported and hence harder to diagnose. CM errors such as syslog messages and alarms should be raised when parity memory errors occur. [PR1059137](#)
- Juniper Networks VSA length above 2K bytes is not supported. Using authorization parameters above this length would result in an incorrect authorization setting for the user. [PR1072356](#)
- When deleting some uncommitted configurations on the active Routing Engine, the rpd process on the backup Routing Engine might restart because of the following error: **Unable to proceed with commit processing due to SIGHUP not received. Restarting to recover**. [PR1075089](#)
- In XM-based multi-LU systems (MX Series platform with MPC3E, MPC4E, MPC5E, MPC6E, NG-MPC3, NG-MPC2, or T4000 with T4000-FPC5-3D linecard), we have multiple LUs representing the same Packet Forwarding Engine complex, we designate the BFD processing to a dedicated LU (LU 0), called as anchor LU, the rest of the LUs (LU 1, LU 2, LU 3) are called non-anchor LUs. When the inline BFD packets punt from non-anchor LU to the anchor LU, it does not have the 'interface-group' populated in the packet context, so the packets might not match by the related filter term. [PR1084586](#)

- On MX Series with MPCs and MICs based platform, the BFD session over IRB interface, with static client, might not come up with VRRP configuration. [PR1085599](#)
- Service chaining of inline software and NAT is not supported. Currently, when we commit configure with software rule and NAT rules under same service-set on SI interface, commit goes fine. [PR1136717](#)
- We support maximum of 1024 software concentrators with inline-6rd. When we configure more than 1024 SCs that is, 1025 SCs, we see commit error message that is not very informative. [PR1153092](#)
- On MX Series routers with MPCs/MICs, the system might try to access a NULL pointer returned from hardware state lookup of the logical interface when the system runs out of memory, which might result in FPC crash. [PR1163606](#)
- The delegated BFD session over aggregated Ethernet interface failed to come up after FEB switchover with FEB redundancy group (1:1 and 1:N). [PR1169018](#)
- With sampling configured, if AS paths change, over a period of time in network, the stale AS paths might be seen in the sampler database of JNH memory. [PR1189689](#)
- Because of a software bug, if an operation related to configuration modification like load merge, load override, commit check, or a commit is done for large configuration files, then mustd crashes and the following error message is seen. **error: Child /usr/sbin/mustd dumped core (0x86)**. [PR1200813](#)
- Several files are copied between Routing Engines during **ffp synchronize** phase of the commit (such as `/var/etc/mobile_aaa_ne.id`, `/var/etc/mobile_aaa_radius.id`). These files are copied even if there was no corresponding change in the configuration, thus unnecessarily increasing commit time. [PR1210986](#)
- IPv6 traffic learned on an L2, bridge, or multilink interface and when it has been traversed through MPLS, core random packets might get classified incorrectly by the fabric, resulting in a packet loss. [PR1223566](#)
- Configuration changes under logical system with LSYS user does not take effect on a single commit with fast-synchronize enabled. [PR1265139](#)
- This issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have XE/GE links toward the downstream switch and LAG bundles as uplinks towards upstream routers. The XE/GE link is part of the physical loop in the topology. Spanning tree protocols such as VSTP, RSTP, and MSTP are used for loop avoidance. Some MAC addresses are not learned on router when LAG bundles that are part of such bridge domains are flapped and other events such as spanning tree root bridge change occur. [PR1275544](#)
- When you see T3 with regard to hardware timestamps, and when you see $T4 < T1$ and $T3 - T2$ because of NTP synchronization, those samples for any measurements (including round-trip time, ingress and egress jitter and delays) are not considered. However, they are reported under error statistics. In these cases, probes are received from the real-time performance monitoring server with correct timestamps and hence they are marked as successful probes. However, they are not used for any measurements. [PR1300049](#)
- When an MPLS unicast next hop gets removed (for example, because of a network convergence event), the statistics associated with that next hop can be erroneously

added to the output statistics of respective physical interface, causing false abrupt increments in output packet and byte count. Logical interface statistics and CoS queues' counters are not affected by this problem and still provide accurate data. [PR1338581](#)

Routing Protocols

- In rare cases, rpd might generate a core file with error **rt_notbest_sanity: Path selection failure on** The core is 'soft', which means there should be no impact to traffic or routing protocols. [PR946415](#)
- On MX Series routers, when an instance type is changed from VPLS to EVPN, and in the same commit an interface is added to the EVPN instance, the newly added EVPN interface might not be able to come up. [PR1016797](#)
- In scaled configurations toggling from 64 bit to 32 bit rpd at the same time that Rosen MVPN routing instances are deleted resulting in a kernel core on the back up Routing Engine. [PR1022847](#)
- The static/static access routes pointing to an unnumbered interface are getting added in the routing table even if the interface is down. In this case, when GRES is disabled, this type of route will never be added in the routing table after Routing Engine switchover. [PR1064331](#)
- When multiple addresses are configured on an interface, if the interface has "interface-type p2p" configured under OSPF and the router does not receive any OSPF packets from one of the IFAs, the OSPF state will not go down for the corresponding adjacency. It should have no impact on route learning, but it might cause confusion for troubleshooting, when peering with Cisco devices, which have multiple addresses configured as secondary addresses. [PR1119685](#)
- A few seconds of traffic loss is seen on some of the flows when PE-CE interface comes up and PE starts learning 70,000 IPv4 prefixes and 400 IPv6 prefixes from CE during L3VPN convergence. [PR1130154](#)
- Generate route doesn't inherit the next hop from the contributing route in an L3VPN case in which the contributing route is learned via MP-BGP. The next hop remains as rejected for the generated route. [PR1149970](#)
- With Shared Risk Link Group (SRLG) enabled under corner conditions, after executing the command "clear isis database", the rpd might crash because the IS-IS database tree gets corrupted. [PR1152940](#)
- Junos marks hidden routes with a negative route preference. The router policy-statement was explicitly setting a preference value for the BGP routes including the hidden routes while the routes are imported into the vrf routing table, overriding the negative one, but not checking for AS_PATH loops within the vrf context. This results in the hidden routes become active in the vrf routing table. [PR1165781](#)
- The rpd might crash when a new PIM interface with the same SNMP index and name as the one that already exists is added to SNMP index tree. The fix is to allow the new PIM interface being added by removing the old one from both the name and SNMP index tree. [PR1178589](#)

- Here are the results when L1 is disabled for Lo0 {master}[edit] user@router# run show isis interface IS-IS interface database: Interface L CirID Level 1 DR Level 2 DR L1/L2 Metric lo0.0 3 0x1 Disabled Passive 0/0 Here are the results when L2 is disabled for Lo0 {master} user@router> show isis interface IS-IS interface database: Interface L CirID Level 1 DR Level 2 DR L1/L2 Metric lo0.0 3 0x1 Passive Disabled 0/0 [PR1202216](#)
- When multiple labels become stale in stale-label-holddown-duration (default 60 secs), it restarts the timer and accumulates all the stale-labels without getting deleted. This might cause memory for allocating labels to be exhausted and then MPLS traffic might be affected due to abnormal/failing label allocation. [PR1211010](#)
- In the context of a large number of configured VPNs, routes changing in the midst of a bgp path-selection configuration change can sometimes lead to an rpd core file. This core file has been seen with the removal of the "always-compare-med" option. [PR1213131](#)
- The rpd might crash if ISIS process tries to send out packet when flow control is not ready. [PR1214947](#)
- On Junos devices with the Graceful Restart, the restarting node might send "End of RIB" maker too soon to its helper nodes, before the actual routes updates are completed, and causing traffic loss. [PR1225868](#)
- The routing protocol process (rpd) on the backup Routing Engine might restart unexpectedly upon the addition of a new L2VPN routing instance. Following a major change in the configuration of an L2V PN routing instance (such as changing instance type on-the-fly, or routing instance rename), the routing instance and all its data structures gets successfully deleted/changed on the master routing-engine. On the backup however, this might not get done, and therefore, the backup might still hold the routing instance with the original name (or original type), in addition to another newly created routing instance reflecting the configuration change. Due to this, one of the routing instance parameters, the "sync-id" is freed on the master Routing Engine, but remains in use on the backup. Later (can be much later), another routing instance gets added in an unrelated provisioning activity. If the sync-id that gets allocated for the new instance is the same as the one for the instance that was not properly deleted in the first change, then rpd on the backup Routing Engine will crash due to an assert, and then restart. [PR1233514](#)
- The distributed Bidirectional Forwarding Detection (BFD) sessions (such as multihop BFD, BFD over AE or MPLS OAM) need to select an anchor Packet Forwarding Engine (PFE). If the connection between Periodic Packet Management Process (ppmd on RE) and Periodic Packet Manager (ppman on PFE) breaks and reconnects due to some reason, the newly chosen anchor FPC might be different than previous one, it is possible that the stale BFD sessions might be seen on previous anchor FPC. [PR1246363](#)
- LDP OSPF are 'in sync' state and the reason observed for this is "IGP interface down" with ldp-synchronization enabled for OSPF. user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down

config holdtime: infinity As per the current analysis, "IGP interface down" is observed as the reason because although LDP notified OSPF that LDP sync was achieved, OSPF was not able to take note of the LDP sync notification, because the OSPF neighbor was not up yet. The issue is under investigation. [PR1256434](#)

- In BGP configuring scenario, the following log entry might be seen in the messages logs under normal operation and should be ignored: rpd[11156]: %DAEMON-3: bgp_rt_send_msg_attr: too big attributes: avail 123 [PR1276758](#)
- If a router works as a graceful restart helper during a peering establishment, the newly established peer might lose some of the negotiated capabilities and might interpret the updates incorrectly. This can cause peer drops or invalid routes to be received. [PR1293174](#)
- Multicast flow interruption might be observed on a transit router in a Protocol Independent Multicast scenario such as the following: (*,G) join is received on one interface, and (*,G) join and (S,G,RPT) prune are received on another interface, which then receives a (*,G) prune. Multicast flows on the first interface get reset and are interrupted for a short period of time (for example, 1 second). [PR1293900](#)
- When BGP family inet labeled-unicast protection is configured, a BGP bypass route might be installed in inet.2. At the same time, if inet.2 is used as the RPF table, the bypass route might be used to perform an RPF check, which will lead to RPF check failure. [PR1310036](#)
- BGP prefixes with three levels of recursion for resolution will get stuck with a stale next hop at the first level after a link-down event, even though the second and the third level routes have updated next hops. [PR1314882](#)
- When two RR (Route-reflector) routers use PIC (protect core) to protect each other's BGP-LU (labeled-unicast) LSP, endless label oscillation might be seen. [PR1318093](#)
- When route target filtering (RTF) is configured for VPN routes and multiple BGP session flaps, there is a slight chance that some of the peers might not be able to receive the VPN routes after the flapped sessions come up. [PR1325481](#)
- With BGP/LDP/ISIS configurations, deleted ISIS routes may still be present in the RIB. The PR does not affect or have any impact on route selection or other functionality of RPD. Just that deleted ISIS routes don't get removed with specific configurations. [PR1329013](#)

Services Applications

- Session counters for clear text traffic are not updated after decryption. Decrypted packet count can, however, be obtained by running the following command **show security group-vpn member ipsec statistics**. [PR1068094](#)
- The issue occurs with multiservices PICs. When you retain SAs in kmd (ipsec-key-management daemon) when the PIC is offline and set them as "not installed", if you reinstall those SAs when the PIC comes online, sometimes reverse routes have reference count problems if the SAs are retained during PIC restart. This issue causes incorrect next hops of the reverse routes. [PR1285907](#)

- ## Subscriber Access Management

- ## User Interface and Configuration

- ## VPNs

- 191

the backup Routing Engine rpd might handle messages and events unsuccessfully and crash. There is no functional impact. [PR1119684](#)

- For a next generation MVPN using ingress replication provider tunnels, if both IPv4 and IPv6 are configured, when receiver PE advertises different labels for IPv4 and IPv6 in type-1 BGP route, the source PE creates two provider tunnels to carry IPv4 and IPv6 traffic both causing duplicated multicast traffic. [PR1128376](#)
- In a next generation MVPN scenario, when **forwarding-cache timeout never non-discard-entry-only** is configured for an MVPN instance, the cache lifetime is shown as forever in the output of the CLI command **show multicast route instance X extensive** and the route disappears after 7-8 minutes. [PR1212061](#)

- See Also**
- [New and Changed Features on page 80](#)
 - [Changes in Behavior and Syntax on page 131](#)
 - [Known Behavior on page 165](#)
 - [Resolved Issues on page 192](#)
 - [Documentation Updates on page 312](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 320](#)
 - [Product Compatibility on page 329](#)

Resolved Issues

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

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

- [Resolved Issues: 15.1R6 on page 192](#)
- [Resolved Issues: 15.1R5 on page 212](#)
- [Resolved Issues: 15.1R4 on page 234](#)
- [Resolved Issues: 15.1R3 on page 253](#)
- [Resolved Issues: 15.1R2 on page 289](#)

Resolved Issues: 15.1R6

- [Class of Service \(CoS\) on page 193](#)
- [Forwarding and Sampling on page 194](#)
- [General Routing on page 194](#)
- [High Availability \(HA\) and Resiliency on page 201](#)
- [Infrastructure on page 201](#)
- [Interfaces and Chassis on page 202](#)
- [Layer 2 Ethernet Services on page 203](#)

- [MPLS on page 204](#)
- [Multicast on page 206](#)
- [Network Management and Monitoring on page 206](#)
- [Platform and Infrastructure on page 206](#)
- [Routing Policy and Firewall Filters on page 209](#)
- [Routing Protocols on page 209](#)
- [Services Applications on page 210](#)
- [Subscriber Access Management on page 211](#)
- [User Interface and Configuration on page 211](#)
- [VPNs on page 211](#)

Class of Service (CoS)

- When the "chained-composite-next-hop" is enabled for Layer 3 VPN routes, MPLS CoS rewrite rules attached to the core-facing interface for "protocol mpls-inet-both-non-vpn" are applied not only to non-VPN traffic (which is the correct behavior) but also to Layer 3 VPN traffic. That is, both MPLS and IP headers in Layer 3 VPN traffic receive CoS rewrite. [PR1062648](#)
- If the hidden command **show class-of-service queue-consumption** is executed many times (in this case, for 100 times), in a rare condition, the cosd process might crash with a core file generated. The core files can be seen by executing the CLI command **show system core-dumps**. [PR1066009](#)
- The **show interfaces queue <if_name>** command has three display options: 1. **show interfaces queue <if_name>** Displays queued/transmitted/dropped packets/bytes for all IFD children. 2. **show interfaces queue <if_name> aggregate** Displays queued/transmitted/dropped packets/bytes for all IFD children except for IFD RTP traffic 3. **show interfaces queue <if_name> remaining** Displays queued/transmitted/dropped packets/bytes for IFD RTP traffic only. Note that unlike queued/transmitted/dropped counters, queues depth values cannot be aggregated. With changes done in this PR, the following is true for queues depth values: 1. **show interfaces queue <if_name>** Displays queues depth values for RTP queues. 2. **show interfaces queue <if_name> aggregate** Displays queues depth values for RTP queues. 3. **show interfaces queue <if_name> remaining** Displays queues depth values for RTP queues. The above logic is the same for physical interfaces, interface-sets, and logical interfaces units. [PR1226558](#)
- On M Series, MX Series, and T Series routers with ingress and egress queueing enabled, input TCP is configured, but no output TCP on the logical interfaces. After you activate or deactivate CoS configuration, the cosd process might crash. [PR1236866](#)
- The following error log message might be seen with Hierarchical CoS and strict-high scheduling configured. **Dec 27 11:08:02.293 mand-re0 fpc1 cos_check_temporal_buffer_status: IFD ge-1/2/1 IFL 358: Delay buffer computation incorrect.^M** If hierarchical scheduler is configured for a physical interface and if guaranteed rate is not set for a logical interface under this physical interface, then the

temporal buffer is configured. The display of error message is valid when guaranteed rate is zero, but it is not valid when guaranteed rate is disabled. [PR1238719](#)

- A round off issue that was leading to a difference in commit behavior of values such as 79m and 79.1m. [PR1252505](#)

Forwarding and Sampling

- On all Junos OS platform, when the ifmon (that is, running the CLI command **monitor interface <interface-name>**) establishes a connection with the Packet Forwarding Engine process (pfed) and runs for a longer duration (as observed, the pfed has been running for more than 11 days), its multiple queries to the pfed might cause the pfed crash due to statistics counter wraparound. [PR1151746](#)
- If a two-color policer is configured on MX Series with MPCs/MICs linecard, more traffic than the limited traffic might be passed when packets size is less than 128 bytes. [PR1207810](#)
- Bandwidth-percent policer does not work on the ps interface, which will result in a commit error. [PR1225977](#)
- Firewall filter family "any" with shared-bandwidth-policer on the MC-AE interface does not reconfigure bandwidth or carve up the policer when standby becomes active after A/S switchover; it drops all packets. [PR1232607](#)
- With a sampling configuration, if you do not define a version for the second flow server, after committing the configuration, the backup Routing Engine might reboot. It might affect routing protocols replicating to the backup Right Engine. [PR1233155](#)

General Routing

- Temp Sensor Fail alarm seen while ASMLC coming up. [PR1036412](#)
- DPD/IKEv2 informational messages are dropped at the peer. Adding vendor ID in the INFORMATIONAL message is causing the peer to drop such packets. [PR1066336](#)
- During IFL clean up "rtsock_peer_unconsumed_obj_add:object already deleted" log messages may indicate that the search bailed citing incorrect results. [PR1085626](#)
- On Junos OS devices, if dot1x is configured, memory leak in kernel might occur that could lead to a system crash. [PR1163782](#)
- In a very rare case, multiple Routing Engine switchovers might result in SNGPMB crash. The SNGPMB is the same thing as Switch Processor Mezzanine Board (SPMB). It is on the line card and contains the LCPU. It also manages locally discovered issues and the switch fabric via the chassis manager thread (CM), which communicates with the fabric manager thread (FM) in chassisd. [PR1176094](#)
- If the MIC-3D-4XGE-XFP is used with MPC2E-3D-NG or MPC3E-3D-NG, the interfaces on the MIC-3D-4XGE-XFP connected to a DWDM device might flap continuously. [PR1180890](#)
- When MS-MIC/MS-MPC is installed on MX, PIC card on MS-MIC/MS-MPC might crash in rare condition. This is a timing issue that might cause traffic loss and has no exact aspect of configurations for triggering that issue. Not as a workaround/restoration, please

refer the external description step 3 for enabling dump under flow-control, which might arise more logs and help for engineer to diagnosis that issue. [PR1182807](#)

- On MX Series platforms, MS-MIC crash might occur. The exact trigger of the issue is unknown; normally, this issue might happen very rarely without any external triggers. The crash might occur with any services configuration, with core files pointing to **Program terminated with signal 4, illegal instruction**. [PR1183828](#)
- On a Junos OS-based platform, CHASSISD_I2CS_READBACK_ERROR error might occur on a single occurrence of I2C read failure. These are transient errors. The errors might be seen randomly without any particular trigger. The fix is to suppress these messages. After the fix, these messages should be seen only when there are three consecutive I2C read failures. [PR1187421](#)
- When VC-Heartbeat is configured, the MX Series Virtual Chassis split detection feature should cause the backup chassis to enter line card isolation mode, powering off all FPCs to force external gear to reroute traffic. A race condition in the mechanism can cause the backup chassis to also become protocol master, and leave its line cards in an operational state, which is undesirable. [PR1187567](#)
- In rare cases, on MX240/MX480/MX960/MX2010/MX2020 Series platforms, MPC4 line card might never come back online after rebooting the chassis by **request system reboot both-routing-engine** command. [PR1190418](#)
- Due to a bug in schema with Junos OS Release 14.1Rx and 15.1rx, administrators will not be able push MPLS configurations to devices that include loose strict tags. [PR1193599](#)
- On an MPC5E, when the "chassis process" (chassisd) notices a high temperature condition on any sensor, it issues a high temperature alarm and increases fan speeds to high. [PR1199447](#)
- An NPC core file might be seen during unified ISSU, and the unified ISSU might fail due to an ISSU ABORT error. [PR1200690](#)
- When Path Computation Element Protocol (PCEP) is enabled and label-switched paths (LSPs) are undergoing changes, like a make-before-break (MBB) change for rerouting, the rpd has to send those updates to the Path Computation Element (PCE). However, when the PCEP session to the PCE goes down, these updates are cancelled, but the rpd fails to completely reclaim the memory allocated for these updates. This causes an increase in the rpd memory every time the connection to PCE goes down while LSPs are simultaneously going through MBB changes. This issue will be especially noticeable when connectivity to PCE goes UP and DOWN continuously. If the connection is in steady state, either UP or DOWN, then the memory leak will not happen. [PR1206324](#)
- The l2ald might thrash when the targeted-broadcast is configured on EVPN irb. [PR1206979](#)
- When using the **show chassis hardware detail** command in Junos OS 15.1 Release or later to display chassis components, the Compact Flash and Hard Disk serial numbers might be truncated to 15 characters. [PR1209181](#)
- The BGP PIC Installs multiple MPLS LSP next hops as Active instead of Standby in Packet Forwarding Engine. This can cause a routing loop. [PR1209907](#)

- When an ARP entry is learned through an AE interface and a route is pointing to that ARP nexthop, the ARP entry will not expire even if the ARP IP is not reachable. This issue occurs due to the route nexthop on the AE interface getting stuck in a unicast state even if the remote end is not reachable, and the RPD is unaware that the ARP is invalid. So, with this resolution, the route nexthop on the AE interface can be shown in the hold state when the remote end is not reachable. [PR1211757](#)
- MS-MPC/MS-MIC might crash when large fragmented (larger than 2048 bytes) traffic goes through an ALG. [PR1214134](#)
- Syslog message : **fpc_pic_process_pic_power_off_config:xxxx:No FPC in slot y** is displayed on empty FPC slots with no PIC power off configured by committing configuration change under chassis hierarchy. [PR1216126](#)
- This issue happens only with RLT configuration and only on Junos OS 16.1 and later releases. [PR1216991](#)
- Suspicious log messages like: **vbf_ifl_bind_change_var_walker:363: ifl.pp.54615 (1073796438): FILTER (28) Bind change notify ran for 276701162891 us** can be observed. [PR1217975](#)
- Kernel crash and router reboot might happen when committing RLT configuration. [PR1218326](#)
- On MX Series platforms, if you are replacing an MQ FPC (MPC Type1, 2, MPC 3D 16x10GE) with an XM one (MPC Type 3,4,5 6. 2E-NG, 3E-NG), all other MQ-based cards might report **FI Cell underflow at the state stage**. It Packets will be dropped. [PR1219444](#)
- On MX Series platforms with enhanced subscriber management, performing a configuration commit that changes any dynamic profile data after the system has booted might result in login and logout connections per second (cps) performance degradation for subscribers using the dynamic profile. [PR1220642](#)
- When **fpc-pfe-liveness-check** is configured, Packet Forwarding Engine liveness detection might incorrectly report a Packet Forwarding Engine failure event under a severe interface congestion situation. [PR1220740](#)
- On MX Series platforms Virtual Chassis partial or complete traffic loss for streams via AE interfaces might be observed in certain scenarios. For example, if vcp ports were de-configured and re-configured again, then two consecutive global GRES switchovers were performed and the MPC hosting AE child links was reloaded, traffic loss would be observed after the MPC boots up due to incorrect programming of AE interface on its Packet Forwarding Engine. [PR1220934](#)
- On MX Series with **pppoe dynamic-profile and service-name-table xx** configured, if configuring the prefix or any interface configuration and after committing, the output of **show pppoe service-name-tables xx** displays as **Service Name Table not found: xx**. [PR1221278](#)
- In the dual Routing Engines scenario with scaled configurations, when events such as daemon restart or Routing Engine switchover occur, the ksyncd process or the backup Routing Engine might crash. This could impact the master Routing Engine in a scaled system because states between the master and the standby are synchronized after a crash. [PR1221913](#)

- After Junos OS Release 15.1, the behavior of storage devices enumeration in kernel level has been changed. Device enumeration in legacy software prior to Junos OS Release 15.1 will show CF and Disk as ad0 and ad1, respectively. Device enumeration after Junos OS Release 15.1 will show CF and Disk as ad1 and ad0 instead in the result of **show chassis hardware**. This might be inconsistent for other results of output, such as **show system boot-messages** and **show log messages**. [PR1222330](#)
- During change of authorization (CoA) requests, there are no changes in schedules. Requests are received successfully, but no changes are sent from the CoS side. [PR1222553](#)
- Due to a defect related to autonegotiation in a Packet Forwarding Engine driver, making any configuration change to interface in MIC "3D 20x 1GE(LAN)-E,SFP" might lead to interface flapping. [PR1222658](#)
- In an enhanced subscriber management environment (**set system services subscriber-management enable**) and when the **remove-when-no-subscribers** configuration statement is configured in auto-configure stanza, when the last subscriber logs out (which triggers the dynamic VLAN IFL removal) and immediately a new subscriber logs in before the IFL is set to inactive, dynamic profile deletion might fail. As a result, subsequent subscriber logins will also fail. [PR1222829](#)
- The "unnumbered-address" under the dynamic profile shows the wrong value. [PR1222975](#)
- The problem of tunnel stream getting misconfigured for LT interfaces was due to internal programming and has been corrected to evaluate multiple LT interfaces for FPC and PIC slot combination. [PR1223087](#)
- In an MX Series platform Virtual Chassis with subscriber management environment, the bbe-smgd process might leak memory in the backup Routing Engine when running continuous subscriber login logout loop tests. It seems memory utilization increases with each login logout loop till it reaches 809 MB and it does not increase beyond that. [PR1223625](#)
- In a PPPoE subscriber scenario, after the demultiplexer underlying interface AEx is changed to AEy, the source MAC used for PPPoE handshake is still the old AEx interface's MAC. This causes PPPoE clients to fail as the PADR packets from the client are dropped due to the MAC address mismatch. [PR1224190](#)
- In a subscriber management environment log message "vbf_ifl_bind_change_var_walker:377: ifl .demux.22698 (1073764522): IFL TCP (38) Bind change notify ran for 1480 us" can often be seen. This log message is generated when time needed to complete execution of the routine exceeds 1ms, it is harmless and can be ignored. However, sometimes time calculation yields incorrect results and this issue has been corrected via this PR. [PR1229967](#)
- The Routing Engine CPU used chassis temperature to decide fan speed. This PR has been fixed to use the real Routing Engine CPU temperature to decide the temperature threshold. [PR1230109](#)
- On all platforms, for IPv6 static routes derived from weighted LSPs, unequal load balance does not work. [PR1230186](#)

- The Random Load Balancing feature does not function; all traffic goes to one of the load-shared egress links instead of being shared across all the links. [PR1230272](#)
- Due to a bug in Junos OS code, the interface statistics remain unchanged post-ISSU on the MXVC platform. This in turn leads to the RADIUS volume accounting value remaining unchanged post ISSU. This is a day-1 issue seen on the MXVC platform only after Junos OS Release 14.1. [PR1230524](#)
- Unsuccessful DCE-RPC ALG sessions result in stale gates and lead to MS-MPC/MS-MIC restart when the gates clean up occurs after timeout. [PR1230868](#)
- The dynamic-profile service filter matches the traffic that is not defined in the prefix-list applied to the filter. As a result, the filter does not work as expected or even match all the traffics. [PR1230997](#)
- The ICMP identifier is not translated back to the expected value during traceroute for TTL exceeded packets on NAT using Multiservice MPC. This occurs for ICMP ID >255 and causes all hops (except first and last) appearing as "*". [PR1231868](#)
- Input framing errors increment on interfaces connected to MPC2E-NG with 4x10G MIC when interface is configured in "wan-phy" mode. [PR1232618](#)
- On the XQ-based linecard, in a rare condition, if offline/online the FPC or link flap, some error messages might be seen. [PR1232686](#)
- High MPC5 CPU on a scaled setup with 64 - 128,000 subscribers due to XQ background service that collects internal statistics. [PR1233452](#)
- When you set port-mirror to the MX Series router, LSP ping might fail and IP packets with options will not get mirrored due to the following unexpected echo reply from DUT: <----- echo request -----> echo reply
[R1]-----[DUT]-----[R2] A | -----> echo reply (unexpected behavior)
| mirror [PR1234006](#)
- After the backup Routing Engine is replaced, the new backup Routing Engine cannot synchronize with the master Routing Engine if **dynamic-profile-options versioning** is configured. This is because the code checks if any dynamic profile is configured before enabling dynamic-profile-options versioning. If so, it throws a commit error. But there is no need to check when the Routing Engine is in backup state. [PR1234453](#)
- KRT queue stuck happening because of socket buffer is sending some junk value to kernel and kernel is returning error 'EINVAL -- Bad parameter in request'. [PR1234579](#)
- When non-Juniper SFP is used in MIC-3D-20GE-SFP-E or MIC-3D-20GE-SFP-EH MIC, the ISR 2 (MIC error interrupt) might be running off over 2.5 second due to unknown reason, and then the FPC host the MIC might be restarted and crashed. The fix add interrupt throttling for MIC interrupt and restarting the MIC if interrupts are more than the threshold (> 2500 per 5min). [PR1235475](#)
- On MX Series platform, when per-packet load sharing is enabled under the aggregated Ethernet interface, egress traffic over the aggregated Ethernet interface might be dropped unexpectedly. [PR1235866](#)

- When PIC-based MPLS J-Flow is configured and MPLS packets are being sampled at the egress (to be sent to the service PIC), the sampled packets do not reach service PIC which results in no MPLS J-Flow flows getting created. [PR1236892](#)
- In an MX Series Virtual Chassis subscriber management environment, LI-enabled DHCP subscribers might experience packet drops because of MAC validation errors in the FPC. This issue was seen only when connecting the subscribers for the first time after rebooting the system. [PR1237519](#)
- DNS server IP addresses are not present in the output of **show subscribers extensive** for DHCP subscribers when DNS configuration is provided from the access-profile or pool. When such data is provided from RADIUS, the output is correct. The issue is cosmetic: DNS addresses are provided to subscribers. [PR1237525](#)
- Due to lack proper boundary checks in code, the MS-MPC might crash when receiving internally corrupted frames from other FPCs that have hardware failure or incorrect rewrite programming. [PR1237667](#)
- Increased support of number of Routing Instances from 4K to 64K. [PR1237854](#)
- MX Series platform is sending accounting interim without an update-interval configuration statement. [PR1239273](#)
- Trace route will not resolve VRF loopback address where system integrator and pseudointerface exist. [PR1240221](#)
- Subscriber Management: MIB ifJnxTable is not supported for subscriber interfaces. [PR1240632](#)
- Session database synchronization might fail if the master Routing Engine or the master chassis in an MX Series Virtual-Chassis configuration (VC-M) is power cycled. [PR1241162](#)
- In some cases, untagged bridged traffic might not be mirrored on the second port of the mirrored group. If untagged bridged traffic is to be mirrored/sent on two different interfaces of the mirrored group, traffic might be mirrored/sent only on one of the mirrored interfaces/ports. [PR1241403](#)
- Routes learned over EBGp multipath peering might not get installed in the forwarding table, resulting in traffic being discarded for the affected destinations. This will only happen if in addition to EBGp multipath there is also a multihop configuration statement enabled for that peering and a unicast reverse path forwarding check is enabled over the involved interfaces. Corresponding routes would end up stuck in the KRT queue and related KRT log messages containing error code **EINVAL -- Bad parameter in request** would be seen in the logs. [PR1241501](#)
- For MX Series Virtual Chassis, some VBF flows are missing after FPC restart. [PR1244832](#)
- The power supply module (PSM) goes to Present State whenever there is a feed failure. The logic is changed to update the PSM state based on the number of feeds connected. [PR1245459](#)
- Fragmented RPC packets can cause the MS-MPC and MS-MIC Service PICs to generate core files when using NAT with the RPC ALGs. [PR1248397](#)
- The bbe-smgd generates a core file when duplicate UID variable names are used for different purposes in the configuration. [PR1248725](#)

- Only one IA-NA dhcpv6 (without a prefix delegation (PD) request) could be bound in case two or more subscribers are provided with the same PD from RADIUS. For example, in case of several customer premises equipment (CPE) devices from a household, all sessions will be provided with the same ACI/ARI. If the username is formed based on ACI/ARI (so the username is the same for all sessions), RADIUS can provide the same PD for all sessions, this will allow only one session to be established even though the CPE devices did not request a PD. [PR1249837](#)
- This issue occurs on MX960 routers with an MPC5E when queues associated with the L4 node get freed but the L4 node itself is not freed. When you try to free the L4 node, because the queues have already been deleted, you will receive a NULL queue code. The MPC crashes with the following message: **qchip_disable_q_rates (q_chip=0x17931598, q_index=73016, q_node=0x0) at ../../../../src/pfe/common/drivers/queue-chip/qchip_rate.c:1801.** [PR1250335](#)
- smihelperd process can crash during subscriber logout process. [PR1250760](#)
- When an IPv6 node receives an ICMPv6 PTB (Packet Too Big) message with MTU < 1280, the node will emit atomic fragments. This behavior might result in denial of service attack. [PR1250832](#)
- On all Junos OS platforms that have rpd, if some interfaces go down, some peers will go down or BGP-RR(route-reflector) re-advertises routes and then the rpd (routing-protocol daemon) process might crash. [PR1250978](#)
- On MX Series with MPC2E-NG/MPC3E-NG, the interfaces of these line cards might not come up when connecting to 3rd party transport switch. [PR1254795](#)
- IRBs that are part of an L3 multicast group allocate ASIC memory when added to the group. A small amount of this memory is not freed when changes are made to the L3 multicast group. This could cause a crash because of an out-of-memory condition if there are continuous changes to multicast groups with IRBs over a long period. [PR1255290](#)
- On some T Series routers, the LSI statistics are not shown in the aggregated Ethernet interface bundles and the input stats counter for the AE interface does not include MPLS traffic. [PR1258003](#)
- MS-MPC/MS-MIC Service PIC constantly generates core files when NAT term calls application-set with no active applications: **application-set EIM_ALG { inactive: application PS3C; inactive: application XBOX1; inactive: application XBOX2; inactive: application XBOX3; inactive: application XBOX4; inactive: application PS3D; inactive: application PS3E; inactive: application PS3F;}**. [PR1258060](#)
- Unable to run **show subscribers extensive** and some other CLI commands after GRES because subscriber-management database is unavailable. The other symptoms of the bug can be messages like **sdb.db: close: Bad file descriptor** and **commit full**. [PR1258238](#)
- In a subscriber service environment, the device control daemon (DCD) might restart unexpectedly during commit process after changes to ATM interface configuration. [PR1258744](#)

- It was observed that an authenticated dynamic VLAN interface is removed with an idle timeout if there are no subscribers on top and if "remove-when-no-subscribers" is configured at the auto-configure stanza. Such dynamic VLAN interface should be removed after its idle timeout expires and if it has stayed idle during this period. [PR1262157](#)
- MX Series use incorrect routing table to send out the ICMP network unreachable message back to the source thus might cause some problem on the end user CPE. [PR1263094](#)
- Dynamic VLAN interface is logged out after reaching idle timeout even though there is a client session (pppoe or dhcp) above it. The proper behavior is to keep the dynamic VLAN interface in case of a client session (pppoe or dhcp) is present above the dynamic VLAN interface. [PR1263131](#)
- It is possible to see a bbe-smgd core under certain boundary conditions on the standby Routing Engine with certain specific configurations. Since the core is on the standby no disruption in service is expected and system recovers from this condition. [PR1267646](#)

High Availability (HA) and Resiliency

- When nonstop routing (NSR) is configured in a group, and that group applied to routing options, NSR sometimes fails. To prevent NSR failure, configure the **nonstop-routing** statement directly at the **[edit routing-instances routing-instance-name routing-options]** hierarchy. [PR1168818](#)
- On all platforms, when running unified ISSU, the connection might be broken between the master Routing Engine and the backup Routing Engine. [PR1234196](#)
- In rare scenario, GRES might not reach ready state and fail to start, due to the fact that Routing Engine does not receive state ack message from Packet Forwarding Engine after performing GRES. This is a timing issue and hard to reproduce. It might also stop Routing Engine resource releasing and then cause resource exhausting. There is no effective method for restoration and the only way is to reboot the system. [PR1236882](#)

Infrastructure

- In an RSVP scenario, when you provision RSVP LSP with ldp-tunneling enabled and these LSPs are configured with link protection, continuous kernel logs and an LDP statistics timeout error might be seen when executing **show ldp traffic-statistics**. [PR1215452](#)
- Polling SNMP QoS queue statistics along with physical interface statistics might result in flat values for QoS queue statistics. The flat values could give a false impression that spikes are happening in the queues. [PR1226781](#)
- On all Junos OS platforms and on the router with PIM enabled that has a local receiver, stale next hops are present because they did not get deleted by daemons due to a timing issue. [PR1250880](#)

Interfaces and Chassis

- In rare conditions, FPC might crash when CLI command **request chassis mic offline fpc-slot <fpc-slot> mic-slot <mic-slot>** or **request chassis pic offline fpc-slot <fpc-slot> pic-slot <pic-slot>** is executed. This is due to a software defect in which SFP diagnostics polling function tries to access already destroyed SFP data structure by MIC/PIC offline. [PR1204485](#)
- The dcd cannot start after router reboot due to nonexistent logical interfaces referenced in **demux-options underlying-interface**. [PR1216811](#)
- In Junos OS Release 14.2 and later, if asymmetric-hold-time, delegate-processing, and preempt hold-time are configured, when the neighbor's interface comes up again, the asymmetric-hold-time feature cannot be used as expected. [PR1219757](#)
- Previously the same IP address could be configured on different logical interfaces from different physical interfaces but in the same routing instance. Only one logical interface was assigned with the identical address after commit. Such behavior could cause confusion: there was no warning during the commit, only syslog messages indicating incorrect configuration. With the fix, it is not allowed to configure the same IP address (the length of the mask does not matter). [PR1221993](#)
- The configuration change where for a static VLAN demux interface the underlying physical interface is changed to one with a lower bandwidth (for example, from xe to ge) can fail with the following error: **error: Bandwidth on IFL demux0.7000 cannot be greater than that of its IFD**. [PR1232598](#)
- On an MX Series platform acting as a broadband network gateway (BNG), in Point-to-Point Protocol (PPP) scenario the router can send LCP Terminate-Ack packet after the PPP over Ethernet (PPPoE) PPPoE Active Discovery Terminate (PADT) packet. This behavior does not follow RFC 2516, which explicitly demands that when a PADT is sent, no further PPP traffic is allowed to be sent using that session, including normal PPP termination packets. [PR1234027](#)
- On M7i and M10i devices, jnxOperatingState shows 1 (unknown) for Fan Tray 1.

```
user@router> show snmp mib walk jnxOperatingDescr | match 4.2
jnxOperatingDescr.4.2.0.0 = Fan Tray 1 jnxOperatingDescr.4.2.1.0 = Fan Tray 1 Fan 1
jnxOperatingDescr.4.2.2.0 = Fan Tray 1 Fan 2 jnxOperatingDescr.4.2.3.0 = Fan Tray 1 Fan 3
jnxOperatingDescr.4.2.4.0 = Fan Tray 1 Fan 4 jnxOperatingDescr.4.2.5.0 = Fan Tray 1 Fan 5
jnxOperatingDescr.4.2.6.0 = Fan Tray 1 Fan 6 jnxOperatingDescr.4.2.7.0 = Fan Tray 1 Fan 7
jnxOperatingDescr.4.2.8.0 = Fan Tray 1 Fan 8 user@router> show snmp mib walk jnxOperatingState | match 4.2
jnxOperatingState.4.2.0.0 = 1 jnxOperatingState.4.2.1.0 = 2
jnxOperatingState.4.2.2.0 = 2 jnxOperatingState.4.2.3.0 = 2 jnxOperatingState.4.2.4.0 = 2
jnxOperatingState.4.2.5.0 = 2 jnxOperatingState.4.2.6.0 = 2 jnxOperatingState.4.2.7.0 = 2
jnxOperatingState.4.2.8.0 = 2 .
```

[PR1237255](#)
- If the maximum transmission units (MTUs) on BNG and CPE sides have different values, the MX Series router might calculate the MTU value for the corresponding pp0 logical interface incorrectly. [PR1240257](#)
- If more than one logical interface (IFL) is configured under the same physical interface (IFD), and VRRP is configured on one IFL without VLAN and the lower unit number IFL has VLAN configuration, VRRP incorrectly carries the VLAN information from the lower

unit number IFL to this logical interface configuration. As a result, VRRP might get stuck (state: unknown, VR State: bringup). This might happen if VRRP is configured on the physical interface with flexible-vlan-tagging or the lt interface without flexible-vlan-tagging. [PR1247050](#)

- When using static demux VLAN interfaces the Link Local address will not be synced between kernel and subscriber management demon. When using router advertisement on static VLAN Demux interface and not in IP dynamic profile, a Router Solicit from customer equipment might not be answered by the MX Series. This is dependant on which address the CPE is using. In this PR, the option to configure the MX Series to use EUI-64 address for the demux VLAN, will ensure that the addresses are synchronized between the demons. [PR1250313](#)
- On Junos OS platforms, the cfmd process runs by default. When bridge-domain is configured, if you commit a configuration related to physical interface/Logical interface (IFD/IFL), cfmd memory leak might occur due to a software defect. As a result, the memory leak can cause cfmd to crash. [PR1255584](#)
- MIC-3D-20GE-SFP-E or MIC-3D-20GE-SFP when reading out periodically SFP diagnostic information, due to misbehaving SFP or noise on the I2C BUS, SFP thread might be hogging and watchdog check will restart the MPC to recover. Enhancements of such error handling will prevent the SFP thread hogging and MPC restart. [PR1260517](#)

Layer 2 Ethernet Services

- When GRES is enabled, after Routing Engine switchover, the local MAC address is not learned anymore from local the CE router in the VPLS instance because of spanning-tree "discarding" in the kernel table. [PR1205373](#)
- Problems with IPv4 HTTP traffic forwarding for dual stacked PPPoE client occur after upgrade from Junos OS Release 14.1X50 to 15.1R4. In this scenario, the user requested two addresses in the DHCPv6 Solicit, an IA_NA and an IA_PD. The server was configured to respond with an IA_PD from a local address pool. The IA_NA was assigned with RA and no address pool for IA_NA was configured at the server. Per RFC, The status codes returned in DHCPv6 Advertise/Reply PDUs from the server when an IA_NA address could not be assigned, should be NO_ADDRS_AVAIL. This was the behavior in Junos OS Release 14.1x5-D150. However, a regression caused this status code to be changed to NO_BINDING instead of NO_ADDRS_AVAIL, in Junos OS 15.1 Release. The CPE in question was likely not interpreting the NO_BINDING status code as a failure from the server to assign an IA_NA address. The status code should respond with Advertise/Reply PDUs with the IA_NA status code of NO_ADDRS_AVAIL. [PR1224212](#)
- During a unified ISSU process, if the first unified ISSU is aborted for some reason, an internal timer will not be cleaned up, and the new lacpd will be forked up. This causes the second unified ISSU in the backup Routing Engine to be aborted in the daemon prepare phase. It will not proceed further. [PR1225523](#)
- MX Series platforms do not include Delegated-IPv6-Prefix in accounting interim. [PR1231665](#)
- This issue can be seen if CPE is initiating DHCPv6-Solicit with IA_NA, IA-PD and Rapid-Commit Option but the MX Series router sends the DHCPv6 Advertise with Rapid

commit flag even though Rapid-Commit statement is not enabled on the MX Series. [PR1235578](#)

- When DPC cards are used and the **set chassis fpc-pfe-liveness-check** configuration statement is configured, some alarms can be seen on the DPC cards (**/var partition is full**) during upgrading from Junos OS Release 15.1F2 - S12 to Junos OS Release 15.1F2 - S13. When trying to downgrade to 15.1F2-S12 the alarm is cleared, and when upgrading to 15.1F2-S13 the alarm is seen again. [PR1237218](#)
- When LACP is configured in fast periodic along with the **fast-hello-issu** configuration statement, LACP might time out if there is any interface commit operation on the peer router during unified ISSU, which causes OSPF adjacency flapping. [PR1240679](#)
- DHCP-Relay option-82 format changes. [PR1253205](#)

MPLS

- On the P2MP LSP transit router with link-protection enabled, if the LSP is the last subLSP, tearing the last subLSP (for example, a RESV tear message is received from downstream router) might crash the routing process (rpd). [PR1036452](#)
- When you have statically configured ingress and transit LSPs, because of a timing issue, the selfID used by the transit LSP might get allocated to the ingress LSP. Ingress static LSPs do not reuse the same selfID across restarts, whereas transit static LSPs try to reuse the selfID. This situation leads to an RPD crash caused by the collision that occurred when the transit LSP tried to reuse the same ID. [PR1084736](#)
- You can configure both **load-balance-label-capability** and **no-load-balance-label-capability** together. This is incorrect and confusing. [PR1126439](#)
- Log messages like **/kernel: %KERN-3: tag_nh_iff_record_delete_iff:404** are cosmetic and were switched on in another PR by mistake. [PR1171947](#)
- When using RSVP-TE protocol to establish LSPs, a make before break (MBB) might not quit and start again when there is a failure on PSB2 (RSVP Path State Block for new LSP) in some cases where PathErr is not seen. For example, for a PSB2 that is already up and there is PathErr processing for it in place already, in this case, no PathErr is seen, owing to local-reversion and a quick flap. As a result, no rerouting happens even if the TE metric cost is raised. This issue has more chances of occurring only when there is non-default optimize switchover delay. [PR1205996](#)
- When dynamic-tunnel is configured but RSVP signaling is disabled, any configuration that affects dynamic-tunnels could cause the rpd process to crash. [PR1213431](#)
- In a scaled environment, when there are many Unicast NHs that are related to the same transport LSP (for example, the same RSVP or LDP label), MPLS traffic statistics collection might take too much CPU time in kernel mode. This can in turn lead to various system impacting events, like scheduler slips of various processes and losing connection toward the backup Routing Engine and FPCs. [PR1214961](#)
- On MX104 Universal Routing Platforms operating with Layer 3 VPN and configured to allow chained composite next hops for devices handling ingress or transit traffic in the network, packets might not be forwarded after they pass through the generic routing encapsulation (GRE) tunnel. This issue is observed on routers operating with Layer 3

VPN that also include the statement **chained-composite-next-hop ingress** at the **[edit routing-options forwarding-table]** hierarchy level. When configured in this manner, the Packet Forwarding Engine cannot push VPN labels for packets. As a result, packets arriving at the next-hop destination cannot be forwarded. [PR1215382](#)

- If the link/node failure that triggered a bypass persists for a long time, and there are LSPs that do not get globally repaired, multiple stale LSP entries are showing down and listed multiple times in the MPLS LSP. [PR1222179](#)
- In a VPLS environment, if you delete the routing instance, in a rare condition, the rpd process might crash, the routing protocols are impacted, and traffic disruption will be seen due to loss of routing information. This is a timing issue and hard to reproduce. [PR1223514](#)
- In impacted Junos OS releases ldp will import metric for all IS-IS routes that have tags without the configuration statement **track-igp-metric**. Junos OS Release 14.1R3, 14.2R1, and later are impacted by this issue. [PR1225592](#)
- Under certain conditions, the entropy label value being generated might not provide a good load-sharing result. [PR1235258](#)
- On MX Series, the rpd might crash when the RSVP bypass undergoes re-optimization and the re-optimized instance encounters failure before it becomes the main instance. The core files could be seen by executing the CLI command **show system core-dumps**.
Stack trace: #0 0x0000000802ad8bd4 in patricia_node_in_tree () #0
0x0000000802ad8bd4 in patricia_node_in_tree () #1 0x00000000009ec3da in tag_pvc_shortwait () #2 0x0000000000a2fe94 in ted_delete_cc_from_link () #3
0x0000000000a3009d in ted2cspf_cleanup () #4 0x0000000000f27d56 in task_job_create_foreground () #5 0x0000000000f289e5 in task_job_bg_dispatch ()
#6 0x0000000000f24d85 in task_scheduler () #7 0x000000000062b9e2 in main ().
[PR1250253](#)
- With non-stop-routing (NSR) and LDP protocol running, a routing protocol process (RPD) on the backup Routing Engine might consume excessive CPU time if it cannot connect to the RPD on the master Routing Engine. [PR1250941](#)

Multicast

- RPD creates an indirect next hop when a multicast route (S,G) needs to be installed when listeners show their interest to S,G traffic. The kernel would then create a composite NH. In this case, this appears to be a P2MP MCNH that gets created. When any member interface is not a Packet Forwarding Engine specific interface (for example, Vt, LSI, IRB, or any other pseudointerfaces), the kernel throws a message indicating that FMBB cannot be supported. These messages are harmless and do not have any impact. [PR1230465](#)

Network Management and Monitoring

- On Junos OS Releases prior to 15.1R6 and 16.1R4, Digital Optical Monitoring (DOM) MIB jnxDomCurrentTable for 1G SFP interface does not return any value. [PR1218134](#)

Platform and Infrastructure

- In a multicast environment, memory leak might be seen on MX Series with MPCs/MICs-based platform after adding, deleting, and changing multicast groups repeatedly. [PR1160909](#)
- If you configure micro BFD on an aggregate interface, using native-vlan and if native-vlan is configured on one of the logical interfaces, then ARP resolution fails for that logical interface. [PR1172229](#)
- On Junos OS platforms with configuration statement **delta-export** enabled, the delta-export database might not get correctly reinitialized upon one of the following conditions: 1. delta-export is enabled for first time (delta-export is enabled in just this commit). 2. load override (delta-export is enabled in the configuration). 3. commit full (delta-export is enabled in the configuration). Because of these conditions, there is a mismatch in databases in further commits. As a result, the configuration on the backup Routing Engine will be corrupted. [PR1199895](#)
- When you check default configurations about groups junos-defaults, there is no information shown. [PR1201380](#)
- Blank firewall logs for IPv6 packets with next-header. [PR1201864](#)
- With 64-bit rpd, if BGP is applied to an export policy with "from protocol", it might cause an error to filter some routes that are not matched with the value from "from protocol". [PR1206511](#)
- On MX Series platforms installed with both DPC/E and MX Series-based MPC, when DPC/E detects a remote destination error toward a MX Series-based MPC Packet Forwarding Engine, unexpected fabric drops happened. [PR1214461](#)
- In large-scale configurations or environment with high rates of churn, the FPC ASIC memory will become "fragmented" over time. It is possible in an extreme case that memory of a particular size will become exhausted and due to the fragmentation, the available memory will not fulfill the pending allocation. [PR1216300](#)
- MX Series with MPCs/MICs-based linecards might crash after firewall filter configuration change is committed. [PR1220185](#)

- When any MPC line card is offlined, it goes offline via all offline flows and connection is cleaned, but in the end of the offline flow, somehow it delays powering off the line card. The chasd powers off the MPC via I2cs that write the respective power registers, but the hardware is not really powering off. As a consequence, since MPC is still powered on but the connection is down, it will try to reconnect, then start to come up automatically within 10 seconds. This issue does not occur all the time. [PR1222071](#)
- Whenever any event (configuration change, login, logout) happens in the system that has to be logged in the accounting server, auditd will be notified about that event. The daemon that notifies auditd about the event writes the event message to a socket and auditd will read from the socket. After reading the message from the socket, auditd will process the event and send the message to the accounting server for logging. A crash occurs when the auditd reads the event message from the socket. Error that is returned while reading is EAGAIN, which means Resource temporarily unavailable. It means currently there is no data to read. When read operation fails, the process is aborted and a core file is generated. [PR1222493](#)
- NTP peers failed to synchronize in symmetric active mode when there is significant downtime of one peer (for example, due to power maintenance or hardware or software upgrades). [PR1222544](#)
- This is a race condition between database creation and database access. It is rarely reproducible. There is no functional impact of the core. [PR1225086](#)
- MAC entry aging is not updated with Source MAC refresh on MPC4E card at slow traffic rate. [PR1230516](#)
- Firewall filter index mapping becomes incorrect after Routing Engine switchover, because the contents of `/var/etc/filters/filter-define.conf` are incorrectly changed after Routing Engine switchover. [PR1230954](#)
- In AI-Scripts (Advanced Insight Scripts) environment, when some special combination of `jcs:printf(...)` and some special characters (such as `\n \t \\`) at the boundary of the buffer, the scripts process might crash and high RPD memory usage is observed. [PR1232418](#)
- The scale-subscriber license count might increase to an invalid license state with L2TP/LTS clients. This is due to the l2tpd daemon not going through proper state transition on L2TP/LTS clients logout. Hence, license count was not getting updated. The fix will ensure license count is updated on logout regardless of whether the daemon goes through proper state transition or not. [PR1233298](#)
- The increase in CPU utilization on the FPCs and MPCs might periodically go as high as 100% as a result of the microcode re-balancing mechanism. UCODE Re-balancing involves instrumenting instructions within the Micro Kernel to gather data. While the PPE is running the UKERN thread in a tight `thread_yield()` loop waiting for a timer to expire, the UKERN scheduler reports the CPU as being 100% utilized. Replaced the tight `thread_yield()` loop with a `timed_semaphore`. Now the semaphore marks the UKERN thread as blocked, the scheduler does not report 100% CPU utilization. [PR1233390](#)
- Login for flow-tap DTCP-over-SSH service fails when SSH key-based authentication is configured for the flow-tap user. [PR1234464](#)

- MX2010/2020 cannot sample multicast traffic when this multicast is copied to multiple interfaces. This behavior is specific to MX2010/MX2020 with MS-MIC. [PR1237164](#)
- FPC and Routing Engine might stuck in high CPU when DDoS SCFD is turned on. [PR1237486](#)
- Starting in Junos OS Release 13.3, the SRX Series cluster needs to run auditd on both nodes. But on MX Series Virtual Chassis Bm and TXP all LCCs also add auditd. Because LCC and VC-BM do not have a route for an accounting server, the following error is generated: ----- 2565 root 1 96 0 3304K 2620K RUN 0:01 0.00% auditd lcc0-re0:
----- 2398 root 1 96 0 3240K 2536K select 0:01 0.00% auditd lcc1-re0:
----- 2791 root 1 96 0 3244K 2544K select 0:01 0.00% auditd %DAEMON-3: auditd[2398]: sendmsg to 10.233.225.78(10.233.225.78).1813 failed: Network is down %DAEMON-3: auditd[2398]: AUDITD_RADIUS_REQ_SEND_ERROR: auditd_rad_send: sendto/sendmsg: Network is down. [PR1238002](#)
- Due to a regression issue, the presence of errors or traps during unified ISSU might result in LU/XL-based FPC crash. [PR1239304](#)
- An FPC crash or only traffic loss might be seen on MPC1E/2E/3E/4E or MPC-3D-16XGE-SFPP during unified ISSU. This issue occurs because counter memory might get corrupted during unified ISSU. It is a timing issue. [PR1241729](#)
- Auditd might crash when RADIUS servers are not reachable and when there are multiple times of Routing Engine switchover. When we try to send RADIUS requests to non-reachable RADIUS servers, we try for maximum number of times. After the maximum number of tries is reached, we close the socket used to send RADIUS requests. After the socket is closed, we are trying to dispatch next message resulting in crash. Auditd will get restarted automatically after it is crashed, so that RADIUS messages if any present in the queue at the time of crash will be lost. After auditd gets restarted, the next event that has to be sent to RADIUS server, will be sent normally. [PR1250525](#)
- On a router with MPC5Es or MPC6Es, if VPLS or bridging features are configured, it is possible that unicast L2 packets with known MAC addresses are flooded instead of being forwarded to the known ports. [PR1255073](#)

Routing Policy and Firewall Filters

- With rib-groups configured for importing routing information to multiple routing tables, unexpected route refresh might happen when committing configuration change. This issue occurs because of a defect in code related to secondary table list handling. [PR1201644](#)

Routing Protocols

- For devices populated with master and backup Routing Engines and configured for nonstop active routing (NSR) and Protocol Independent Multicast (PIM) configuration, the routing protocol process (RPD) might crash on the backup Routing Engine due to a memory leak. This leak occurs when the backup Routing Engine handling mirror updates about PIM received from the master Routing Engine deletes information about a PIM session from its database. But due to a software defect, a leak of 2 memory blocks (8 or 16 bytes) might occur for every PIM leave. If the memory is exhausted, the rpd might crash on the backup Routing Engine. There is no impact seen on the master Routing Engine when the rpd crashes on the backup Routing Engine. Use the **show system processes extensive** command to check the memory. [PR1155778](#)
- After Routing Engine switchover, a race condition could result in a RIB not registering for route flash. As a result, there might be stale entries seen when routes are withdrawn. This is a rare condition. [PR1170572](#)
- When you have LSPs as OSPF neighbors and run **run show snmp mib walk decimal 1.3.6.1.2.1.14** you get the message request failed, and the command does not complete. [PR1177315](#)
- In a BGP scenario with inet-mdt family configured under protocols BGP, route table <NAME>.mdt.0 might get deleted if it has no routes. As a result, rpd might crash on the backup Routing Engine, and BGP sessions might flap on master Routing Engine. [PR1207988](#)
- When changing the route distinguisher (RD) for an existing routing instance with established MSDP sessions or deleting or deactivating an MSDP session in the configuration, the rpd process might crash, which leads to traffic disruption. [PR1216078](#)
- rLFA OSPF protection path/next hop keeps flapping every 4-5 seconds when metric order policy is configured. [PR1220343](#)
- When the first multicast packet gets fragmented because of its large size, the receiver in an MVPN scenario does not receive all fragments. The fix of this PR will ensure that the software waits until the last fragment of the PIM register packet is received at the rendezvous point (RP) before processing the PIM resolve request. After the last fragment of register packet is received, PIM register state is created and PIM resolve request is triggered to install the multicast route. Hence, all fragments of the register packet will get forwarded to the receiver. [PR1229398](#)
- Junos OS Release 15.1 and later might be impacted by the receipt of a crafted BGP UPDATE which can lead to an routing process daemon (rpd) crash and restart. Repeated crashes of the rpd daemon can result in an extended denial of service condition. Refer to JSA10778 for more information. [PR1229868](#)

- Remote LFA protection might not work for the OSPF route when there is no ECMP to act as a candidate for the PQ node (the PQ node's router ID belongs to a different area). [PR1230322](#)
- In a rare condition after BGP session flaps, BGP updates might not be sent completely, resulting in BGP routes shown in the advertising-protocol table in the local end but not shown in the receive-protocol table in the remote end. [PR1231707](#)
- The routing protocol process (rpd) sometimes is interrupted and halted when it tries to free a session reference block. This can occur when the memory redzone check fails when attempting to free reference memory block. The fail is caused when the redzone check receives an address that is not the beginning of a memory block. [PR1232742](#)
- When MX Series router is running protocol BGP, and policy configuration is modified, an assertion condition might be hit where the routing protocol process (rpd) generates a core file. [PR1239990](#)
- Session uptime in **show bfd session detail** output omits seconds if uptime is longer than 24 hours, which is different from similar output for Label Distribution Protocol (LDP), Open Shortest Path First (OSPF), or Border Gateway Protocol (BGP). Seconds are always included in corresponding outputs for these protocols. [PR1245105](#)
- On all platforms, if multi-area rLFA along with policy is configured, Open Shortest Path First (OSPF) nexthop might keep flapping. [PR1248746](#)

Services Applications

- In an L2TP scenario, when the LNS is flooded by high rate L2TP messages from LAC, the CPU on the Routing Engine might become too busy to bring up new sessions. [PR990081](#)
- IDP policy is trashing with following log messages: **Aug 23 20:56:25 esst480a jlaunchd: idp-policy (PID 8562) started Aug 23 20:56:25 esst480a jlaunchd: idp-policy (PID 8562) exited with status=0 Normal Exit Aug 23 20:56:25 esst480a jlaunchd: idp-policy (PID 8564) started Aug 23 20:56:30 esst480a jlaunchd: idp-policy (PID 8564) exited with status=0 Normal Exit Aug 23 20:56:30 esst480a jlaunchd: idp-policy (PID 8570) started Aug 23 20:56:35 esst480a jlaunchd: idp-policy (PID 8570) exited with status=0 Normal Exit Aug 23 20:56:35 esst480a jlaunchd: idp-policy (PID 8574) started Aug 23 20:56:40 esst480a jlaunchd: idp-policy (PID 8574) exited with status=0 Normal Exit.** On MX Series platforms, running IDP process is stopped from Junos OS Release 14.2 and later. So **idp-policy** configuration does not work. This is an expected behavior. Starting with the fixed versions we have completely deprecated the IDP related CLIs on MX Series platforms and we should not see any IDP related messages. [PR1209351](#)
- The kmd process might monopolize the CPU when continuous polling for IKE-related data through SNMP. This issue is specific to IKE-related SNMP polling and not seen when continuously polling IPsec-related data through SNMP. [PR1209406](#)
- In case of massive flapping of subscribers on M120 platform, a memory leak on IQ2E PIC can happen and it can result in the inability to attach a configured CoS policer to the newly connected l2tp subscriber. [PR1210976](#)

- When loading or rolling back a configuration that removes a service-set and changes where the MS interfaces are assigned, traffic may be blackholed to a series of the existing service-set might be dropped. [PR1223302](#)
- When the stateful firewall flows time out repeatedly, there can be performance degradation on the MS-DPC PIC. This will eventually lead to the MS-DPC being unable to scale to the peak flows that are allowed. [PR1242556](#)
- On Layer 2 Tunneling Protocol (L2TP) network server (LNS) router L2TP tunnels might be stuck in "Terminating" state after execution of particular sequence of CLI commands. Deactivation of tunnel-group on LNS leads to cleanup of all logged-in L2TP subscribers and L2TP tunnels. If the **clear services l2tp tunnel** command is issued when the cleanup has not been completed, it is possible that the tunnel will not be cleaned up properly and will get stuck in "Terminating" state. [PR1249768](#)

Subscriber Access Management

- In a subscriber management environment with two or more RADIUS servers connected to an MX Series router, syslog is not generated when radius server is mark dead. [PR1207904](#)
- On MX Series routers with dual Routing Engines, after router GRES, if you add a traceoptions filter before GRES is fully completed, the authd process might crash. [PR1234395](#)
- The command **show network-access aaa statistics radius detail** can display incorrect number of messages to RADIUS server in case configured RADIUS server's are continuously flapping. [PR1267307](#)

User Interface and Configuration

- This issue is specific to a router running a Junos OS Release up to 15.1R<x>, which also has authentication-key-chains configured. When the secret for a key is not configured, commit fails with the message **error: configuration check-out failed: daemon file propagation failed**. This issue is not applicable to Junos OS Release 15.1F, 16.1, and later. [PR1213165](#)
- Some configuration objects are not properly handled by "delta-export" (dexp). This leads to an omission of the section of the configuration. [PR1245187](#)

VPNs

- For NG-MVPN, the traffic threshold is ignored if it is configured in a configuration group, then applied to an MVPN instance. If the traffic threshold is configured directly under the MVPN instance, the issue is not seen. [PR1191002](#)
- On Junos OS platforms, only VPLS supports automatic-site-id. Configuring automatic-site-id under the L2VPN instance could cause a rpd core file. The fix has now been provided to add a commit check to disallow configuring automatic-site-id under a L2VPN instance. With this fix, a commit error will be generated if you attempt to configure automatic-site-id under an L2VPN instance. [PR1214328](#)
- The routing protocol process (rpd) might eventually become exhausted and crash when Layer 2 Circuit, Layer 2 VPN, or virtual private LAN service (VPLS) configurations

are committed. These commit activities might create a small memory leak of 84 bytes in the rpd. [PR1220363](#)

- In an NG-MVPN scenario with the `asm-override-ssm` configuration statement for source-specific multicast (SSM) group, if you issue the `clear pim join` command on the source PE router, downstream interfaces get pruned, causing the multicast flow to stop. [PR1232623](#)
- On M Series and MX Series platforms, the L2circuit does not switch from primary to backup and vice versa based on the APS status change, because when APS switchover happens, the PW switchover does not switch to the new APS active neighbor. [PR1239381](#)
- With NSR enabled and a Layer 2 circuit configured, an rpd crash might be observed on the backup Routing Engine when you change the Layer 2 circuit neighbor and then commit the changes. [PR1241801](#)

Resolved Issues: 15.1R5

- [Class of Service \(CoS\) on page 212](#)
- [Forwarding and Sampling on page 213](#)
- [General Routing on page 213](#)
- [High Availability \(HA\) and Resiliency on page 222](#)
- [Infrastructure on page 222](#)
- [Interfaces and Chassis on page 222](#)
- [Layer 2 Ethernet Services on page 223](#)
- [Multiprotocol Label Switching \(MPLS\) on page 224](#)
- [Network Management and Monitoring on page 225](#)
- [Platform and Infrastructure on page 225](#)
- [Routing Policy and Firewall Filters on page 228](#)
- [Routing Protocols on page 229](#)
- [Services Applications on page 231](#)
- [Subscriber Access Management on page 232](#)
- [User Interface and Configuration on page 233](#)
- [VPNs on page 233](#)

Class of Service (CoS)

- In rare cases, after polling "show snmp mib walk jnxCosQstatTxdBytes", cosd coredump might occur due to memory corruption on Junos platform with COS enabled. [PR1199687](#)
- The actual problem seen is Logical Interfaces binded to Routing-instance classifier is not seen under classifier Index inside CFEB. The cause for this Issue was "missing else statement was leading to data getting overwritten for LSI scenario". The same has been Corrected. [PR1200785](#)

Forwarding and Sampling

- The dfwc (daemon that performs as firewall compiler) might fail to get filter information from the kernel in COMMIT_CHECK (config validation) mode. As a result, the filter index is regenerated starting from index 1. This will create the mismatch of filter index as compared to the existing filters in the system. The fix provided will identify and recover the issue. [PR1107139](#)
- Commit gives error as follows when apply-groups is configured under bridge domain. error: Check-out failed for Firewall process (/usr/sbin/dfwd) without details. [PR1166537](#)
- SRRD(Sampling Route-Record Daemon) process doesn't delete routes when the DELETE is received from RPD in few configuration cases. This results in build-up of memory in SRRD daemon and once SRRD reaches the limit, it crashes and restarts itself. This happens only when one certain family is not configured on all of the FPC clients (e.g., FPC with inline J-Flow enabled or PIC with PIC-based sampling enabled is one client). For example, only IPv4 family is configured in all the clients, and IPv6 and MPLS families are not configured for sampling in any of the clients. [PR1180158](#)
- The changes to srrd (sampling route reflector daemon - new architecture for sampling) process between 14.2R5.8 and 14.2R6.5 severely reduce MX80 series available memory and therefore RIB/FIB scaling. [PR1187721](#)
- Starting with Junos Release 14.2R1, FPC offline could trigger Sampling Route Record (SRRD) daemon restart. [PR1191010](#)
- On MX platform with "Enhanced Subscriber Management" mode, if default forwarding-classes are referenced by subscriber filters, commit configuration changes after GRES will be failed. [PR1214040](#)

General Routing

- In the scenario when one interfaces having same IP addresses with a RSVP strict path en-routed interface IP address (for example, subscribed interface borrows the loopback interface IP address scenario, or where one of PE-CE interface inside a VPN instance has the same IP address of the router's uplink RSVP interface in master instance), RSVP-TE would send PathErr to ingress router due to matching to wrong interface which is not RSVP interface but having same IP address with the RSVP interface when checking the explicit route object (ERO). [PR 1031513](#)
- On dual Routing Engine platform with GRES and NSR enabled, after Routing Engine switchover, the rpd might crash when trying to destroy a CNH NH (composite next hop, for example, it would be created in PIM, L3VPN, MVPN scenario and so on) with valid reference on it. It is because that during switchover (while backup rpd switches to master), there is a transition period where rpd switched to master mode but KRT is still in backup mode. If KRT (still in backup mode) receives a CNH addition followed by Route additions using this CNH during this phase, it would result in CNH in KRT with valid route references yet on expiry queue. It is hard to reproduce, in this case, it occurs after Routing Engine switchovers consecutively at two times. [PR1086019](#)
- The configuration support for enabling ingress and egress layer2-overhead is available in dynamic-profile but the functionality is not supported in 15.1R3 and 15.1R4. For example, set interfaces ge-4/2/9 unit 0 account-layer2-overhead ingress 30 set

interfaces ge-4/2/9 unit 0 account-layer2-overhead egress 30 With the above configuration, the number of layer2-overhead bytes (30) are not added to the input bytes in traffic statistics. [PR1096323](#)

- During NSR Routing Engine switchover, there might be a control plane black window for inline BFD causing the BFD session to flap. This is a day-1 issue, and tuning the retrans timer would solve the problem. But since these timers have to meet RFC compliance, we cannot really do that. Today we have the retrans timer as 1000 milliseconds. The workaround would be to configure a higher retrans timer value. [PR1105980](#)
- The rpd fails to respond any new CLI routing commands (for example, show mpls lsp terse). Rpd is forking a child process while rpd is processing a show command. When the subprocess tried to exit, it tried to close the management socket being used by the show command. This failure might cause rpd subprocess to crash and generate a core file. It also removes the rpd pid file which prevent rpd from processing any new CLI commands even though original rpd process continues to run normally. [PR1111526](#)
- During initial ramp up of an IPSec session, a race condition might cause the mspmand process crash in rare circumstances. [PR1116487](#)
- On MX Series platform with MS-MPC/MS-MIC in use, due to some reason if the NAT session is freed/removed but without removing timer wheel entry, then it might cause MS-MPC/MS-MIC crash. It is a timing issue where just before invoking the timer wheel callback the NAT session extension got freed/removed. [PR1117662](#)
- On MX Series platforms, the MS-MPC crash might occur. The exact trigger of the issue is unknown; normally, this issue might happen over long hours (e.g. within a week) of traffic run (e.g. running HTTP/HTTPS/DNS/RTSP/TFTP/FTP traffic profile). Core dumps might point to - Program terminated with signal 4, Illegal instruction [PR1124466](#)
- The jsscd might crash in static-subscribers scaling environment (e.g. 112K total subscribers, 77K dhcp subscribers, 3K static-subscribers, 32K dynamic vlans), when this issue occurs the subscribers might be lost. abc@abc_RE0> show system core-dumps -rw-rw---- 1 root field 8088852 Jan 11:11 /var/tmp/jsscd.core-tarball.0.tgz [PR1133780](#)
- In a multicast virtual private network (MVPN) scenario during route churn, the rpd process might crash due to inconsistency multicast next-hop between rpd and kernel. [PR1138366](#)
- On MX Series platforms with FPC3, the octets of IPv4 source and destination addresses in the firewall log are listed reverse; this might affect troubleshooting. The IPv6 log works fine. This is a minor issue, there is no other service impact. [PR1141495](#)
- During route flaps such as (interface flaps or network instability) the Packet Forwarding Engine may reboot or Packet Forwarding Engine may notice next-hop corruption. [PR1151844](#)
- If any linecard crashes early during ISSU warmboot, the CLI might report ISSU success, resulting in a "silent ISSU failure". [PR1154638](#)

- In sampling feature, certain scenarios force handling of the sampled packet at the interrupt context, which may have chance to corrupt the BMEB packet context, and lead to BMEB FDB corruption. [PR1156464](#)
- On MX Series with MPCs/MICs platforms with MPC2-NG/MPC3-NG/MPC3/MPC4/MPC5/MPC6 installed, in rare cases, a very rare hardware error - TSTATE Parity error might occur. It can cause FPC getting stuck, but it will not trigger the error-reporting infra (CMERROR). Fixes have now been provided. [PR1156491](#)
- The default (per-packet load balancing) PPLB export policy created for Ethernet VPN (EVPN) has been removed from JUNOS. It was used to enable per packet load-balance for EVPN routes on certain MX platforms and not all. Now per-packet load balance needs to be configured explicitly. [PR1162433](#)
- On Junos 15.1 and above, after Routing Engine switchover and both Routing Engine reboot, krt queue might get stuck. It's because: under this scenario, agentd creates it's table before rpd reading tables. But after rpd restarting and rebuilding tables, it could not filter an agentd's table out. It might cause slow route convergence or traffic loss. This issue would disappear automatically in 30 minutes. [PR1162592](#)
- On MX Series router with services PIC (MS-DPC/MS-MPC/MS-MIC), the ICMP time exceeded error packet is not generated on an IPsec router on the decap side. [PR1163472](#)
- When the MS-MIC or MS-MPC installed in MX Series router is processing traffic, and the IPsec policy configuration is changed by means of adding or upating a policy, mspmand process crash might occur. [PR1166642](#)
- Sampled continues logging events in trace option file after trace option for sampled deactivated. This can be hit if there is no configuration under 'forwarding-options sampling' but other configuration for sampled is present (e.g. port-mirroring). [PR1168666](#)
- When MS-MPC is used, if any bridging domain related configuration exists (e.g. "family bridge", "vlan-bridge", "family evpn", etc), in some cases, continuous MS-MPC crash hence traffic loss might occur. [PR1169508](#)
- When using Periodic Packet Management process (PPMD, responsible for periodic transmission of packets on behalf of its various clients) related protocols (e.g. LFM, CFM, LACP, BFD, etc), during fabric or SIB online process, possibly, the client session (who establish adjacencies with PPMD to receive/send periodic packets on those adjacencies, such as LFM, CFM, LACP, etc) of PPMD may flap due to CPU hog issue. [PR1174043](#)
- On Virtual Tunnel (vt) tunnel environment with forwarding-class, customer is using AE interface to terminate subscribers on the box and the AE interface has members on two different FPCs, due to a software defect, the mirrored traffic is not going to the correct forwarding class as expected. The issue is also seen when terminate Subscribers and vt tunnel hosted interface are on two different FPCs (Non-AE case). [PR1174257](#)
- When using MS-MPC or MS-MIC service cards, a single pool cannot be used in different service-sets. Separate pools with different names would then need to be used. Additionally, pools created automatically by a source-prefix or destination-prefix statement will not work if the same source-prefix or destination-prefix statement appears in a different service-set. [PR1175664](#)

- MTU discovery may not be working due to lack of VRF info on egress card for BBE Subscriber traffic. [PR1177381](#)
- This is a display issue and doesn't affect functionality of the power, fixing has been added to commands 'show chassis power' and 'show chassis environment pem', when one of the DC PEM circuit breaker tripped. [PR1177536](#)
- CGNAT-NAT64: Few port leak are observed for the EIM/EIF IPv4 traffic(2M sessions) from public side. [PR1177679](#)
- destination-prefix-list support list added for NAT rule with twice-napt-44 translation. Customer will be able to define a prefix list and match it in the NAT rule while using twice-napt-44. [PR1177732](#)
- If "router-advertisement" protocol is configured in client ppp profile, unsolicited RA might be sent before the IPv6CP Configuration ACK is received. [PR1179066](#)
- After One side PE Junos upgrade from the release before 15.1R1 to the release after 15.1R1, due to the construction of es-import-target changed, type 4 routes are not imported and missed in table __default__evpn__evpn.0, which caused both PEs thought itself is DF router and forwarding BUM frames. This will prevent to upgrade Junos in production network. [PR1179443](#)
- On T-series platforms with 10x10GE Type 4 PIC installed, if an interface in such PIC is configured with WAN PHY mode, the CoS configuration on the port will be incorrectly programmed and it might result in unexpected packet drop. [PR1179556](#)
- On dual Routing Engine platforms, if interface changes occur on Aggregate Ethernet (AE) which result in marking ARP routes as down on the AE (e.g. bringing down one of the member links), due to interface state pending operation issue on backup Routing Engine, in race condition, the backup Routing Engine may crash and reboot with an error message (panic:rnh_index_alloc: nhindex XXX could not be allocated err=X). [PR1179732](#)
- In the CGNAT CLI show service alg conversations fails to display parent session status for ALG conversations. [PR1181140](#)
- In case of point to point interfaces and unnumbered interfaces rpd crash might be seen in corner cases on configuration changes. There is potential fix given through this PR to avoid the crash. [PR1181332](#)
- When "dynamic-tunnels" is configured with configuration statement "gre", performing Routing Engine switchover might result in rpd crash. [PR1181986](#)
- Fragmented ALG control traffic is not supported on the MS-MPC. [PR1182910](#)
- With NAT translation-type as napt-44, a few sessions are getting stuck upon deactivating/activating service-set or corresponding applications at a few times with traffic running. The same symptom is seen upon deactivating/activating service-set with traffic running and with 'deterministic-napt44' translation type as well. [PR1183193](#)
- CGNAT Pool stats for "Available address" is shown incorrect for destination pool. Available address shown zero even though destination nat IPs are available [PR1183538](#)
- With BGP add-path and consistent-hash enabled, when a BGP learnt route prefix with multiple paths(next-hop) is installed in the forwarding-table, all the next-hops should

be reachable/resolvable at the time of installing the route in the forwarding-table. However, there might be a chance that any of the next-hops are not resolvable at that time, which will lead Packet Forwarding Engine's incorrect route programming. In this case, traffic forwarded to this prefix will be affected. [PR1184504](#)

- When IPv4 firewall filter have 2625/32 destination in prefix-list, filter attached to subscriber interface is found broken. [PR1184543](#)
- Starting with 15.1F5, the splitting of destination NAT pools across AMS members will be prevented. Currently with AMS interfaces, dnat44 pools do not get split. However, all twice-NAT destination pools are split. This is not needed and this change makes it so (source pools are split or/and hashing is based on source so there is never any chance of conflict). Please work with Francois to get details. [PR1184749](#)
- Continuous reporting of the following messages might be noticed sometimes while bringing up all IFD/IFL/IFF states at once.

```
Apr 1 11:16:05 mx2020-1 dot1xd[16641]: %--: task_receive_packet_internal: knl Ifstate packet from zero-len socket 8 truncated. Apr 1 11:16:05 mx2020-1 dot1xd[16641]: %--: Free allocated bufp:(a433004) buflen:(16384)task_receive_packet_internal: knl Ifstate packet from zero-len socket 8 truncated. Apr 1 11:16:05 mx2020-1 dot1xd[16641]: %--: task_receive_packet_internal: knl Ifstate packet from zero-len socket 8 truncated. Apr 1 11:16:05 mx2020-1 dot1xd[16641]: %--: task_receive_packet_internal: knl Ifstate packet from zero-len socket 8 truncated. Apr 1 11:16:05 mx2020-1 dot1xd[16641]: %--: Free allocated bufp:(a433004) buflen:(16384)task_receive_packet_internal: knl Ifstate packet from zero-len socket 8 truncated. Apr 1 11:16:05 mx2020-1 dot1xd[16641]: %--: task_receive_packet_internal: knl Ifstate packet from zero-len socket 8 truncated.
```

During syncing of ifstate dot1xd, try to read all the Physical Interfaces/ifl/iff state at once. In scale scenario, the size of these information will be very high. It may exceed demon rlimit / memory availability.[PR1184948](#)

- In IPv6 environment, adding a link local neighbour entry on subscriber interface then adding a new lo0 address, if delete this neighbour entry and the subscriber interface, due to software defect, the nexthop info is not cleaned properly, the rpd process might crash. The routing protocols are impacted and traffic disruption will be seen due to loss of routing information. [PR1185482](#)
- When ams-interface is configured in warm-standby mode without adding any members, configuration commit will lead to rdd core. [PR1185702](#)
- AMS redundant interfaces not listed under possible-completions of operational commands. [PR1185710](#)
- In IPv6 environment with graceful Routing Engine switchover (GRES) enabled, when a new prefix (global address) is added on the donor interface (in this case, loopback interface), and then perform GRES, the ksyncd process crash might be observed due to kernel replication error. [PR1186317](#)
- When both AMS-redundant interface and AMS-load-balancing interface is configured in the system, 'Not a deterministic nat pool' syslog is generated whenever deterministic-nat show cli command 'show services nat deterministic-nat nat-port-block' is executed. [PR1186723](#)

- JUNOS might improperly bind Packet Forwarding Engine ukernel application sockets after ISSU due to a bug in IP->TNP fallback logic. Because of that bug, threads running on the ukernel that relay on UDP sockets can experience connectivity issues with host, which in turn can lead to various problems. For instance, sntp (simple network time protocol) client might fail to synchronize time, which in turn might lead to other problems such as failure in adjacency formation for HMAC authenticated protocols. [PR1188087](#)
- By default SNMP will cache SNMP values for 5 seconds. Sometimes kernel will cache these values for longer duration. This PR will correct the caching behavior. [PR1188116](#)
- The command "request system reboot both-routing-engines local" on VC-Mm will reboot only one Routing Engine on an MX-VC, with this fix, it will reboot both Routing Engines of local chassis. In addition, this fix also removes the "set virtual-chassis member <n> role line-card" configuration option on an MX-VC because this option is not supported on MX-VC as designed. [PR1188383](#)
- On MX routers, a vulnerability in IPv6 processing has been discovered that may allow a specially crafted IPv6 Neighbor Discovery (ND) packet to be accepted by the router rather than discarded. The crafted packet, destined to the router, will then be processed by the Routing Engine (RE). A malicious network-based packet flood, sourced from beyond the local broadcast domain, can cause the Routing Engine CPU to spike, or cause the DDoS protection ARP protocol group policer to engage. When this happens, the DDoS policer may start dropping legitimate IPv6 neighbors as legitimate ND times out. Refer to JSA10749 for more information. [PR1188939](#)
- Ingress queuing configuration on MPC2ENG is leading to host loopback wedge due to some bug in the code specific to MPC2ENG; there is a mis-programming in the Junos code for the lookup chip for this type of card. [PR1189800](#)
- When polling an si-interface hosted on an NG-MPC Non-HQoS line card (MPC2E-3D-NG, MPC3E-3D-NG), there always has a 10 sec delay, which might break SNMP polling. [PR1192080](#)
- As described in RFC7130, when LACP is used and considers the member link to be ready to forward traffic, the member link MUST NOT be used by the load balancer until all the micro-BFD sessions of the particular member link are in Up state. [PR1192161](#)
- If a message received from LLDP neighbor contains "Port Id" TLV which has "Interface alias" subtype and is longer than 34 bytes, subsequent running of "show lldp neighbors" might lead to l2cpd crash. [PR1192871](#)
- On MX series with MPC3/MPC4/MPC5/MPC6, the VSC8248 firmware on the MPC crashes occasionally. This PR enhances the existing VSC8248 PHY firmware crash detection and recovery, helping recover from a few corner cases where the existing JUNOS workaround does not work. [PR1192914](#)
- When MoFRR activated, multicast source route flapping leads to corresponding multicast traffic 100% drop. [PR1194730](#)
- On Junos OS Release 15.1R3 and later with Tomcat model BBE release, if a subscriber login/logout which using multicast service, then another subscriber login and also use multicast service, this may cause bbe-smgd core on backup Routing Engine. [PR1195504](#)

- In inline BFD or distributed BFD (in Packet Forwarding Engine) scenario, Packet Forwarding Engine fast reroute is not invoked anymore if the remote peer signals BFD ADMINDOWN message to local node and convergence time is performed based on protocol signaling. [PR1196243](#)
- On platforms running Junos OS with FreeBSD10, if tracing is enabled, due to the log file pointer not being handled correctly for log file rotation, the rpd process might crash when the log file rotates. [PR1196318](#)
- Distributed BFD session using inline-redirection on MX-VC might not work if the ANCHOR Packet Forwarding Engine is not within the same chassis member as the interface where the BFD packet is received from peer device [PR1197634](#)
- L2VPNs or L2Circuit services along with lengthy interfaces descriptions might lead to memory leak in variable-sized malloc block, which in turn results in RPD crash due to "out of memory". [PR1198165](#)
- Problem: ===== The following continuous error messages are generated during 2X100GE CFP2 OTN MIC online on MX2K. This error message means PCI control signal communication failure between Packet Forwarding Engine on MPC6E and PMC Sierra OTN framer (pm544x) on MIC 2X100GE CFP2 OTN. *** messages *** Jul 25 17:39:04.807 2016 MX2K : %PFE-3: fpc0 cmic_pm544x_hires_periodic: error getting counters Jul 25 17:39:04.893 2016 MX2K : %PFE-3: fpc0 Failed in function pm544x_manage_link:2616 Jul 25 17:39:05.267 2016 MX2K : %PFE-3: fpc0 Failed in function pm544x_link_status:2449 Jul 25 17:39:05.267 2016 MX2K : %PFE-3: fpc0 cmic_pm544x_hires_periodic: error getting counters Jul 25 17:39:05.267 2016 MX2K : %PFE-3: fpc0 Failed in function pm544x_manage_link:2616 Jul 25 17:39:05.267 2016 MX2K : %PFE-3: fpc0 Failed in function pm544x_link_status:2449 Jul 25 17:39:05.321 2016 MX2K : %PFE-3: fpc0 cmic_pm544x_hires_periodic: error getting counters Jul 25 17:39:05.408 2016 MX2K : %PFE-3: fpc0 Failed in function pm544x_manage_link:2616 Jul 25 17:39:05.486 2016 MX2K : %PFE-3: fpc0 Failed in function pm544x_link_status:2449 Root cause: ===== Bug was in converting the 32bit PCI shared address to 64 bit address. When the MSB of the 32bit address was set, the conversion was buggy as it type caused it to signed long int, which resulted in extending the sign bit to first 32 bits of the converted 64bit address. The first 32bit of the converted address is expected to be zero as our memory is only 32 bit addressable. Problem appearance on customer deployments:
===== 1. Issue will be seen only when there are large number of nexthops in the Packet Forwarding Engine due to Packet Forwarding Engine anchor feature before the MIC is made online. 2. If the MIC came online without hitting this issue, then there is no chance of hitting this issue later. Because the bug was in the PCI shared memory allocation, which happens only during the MIC online. 3. This issue started showing after the Packet Forwarding Engine anchoring feature, which delayed the MIC online until the next-hops are sync to Packet Forwarding Engine. As a result the MIC is coming online very late and the shared memory allocation is coming from the higher RAM address, which the PMC vendor code porting layer is failing to handle. After the fix from this PR, we should not hit this issue. [PR1198295](#)
- With MPC-NG or MPC5E hardware, the range of the queue weights on an interface is from 0 to 124. As every queue has to have an integer value of queue weight, it might be impossible to assign the weights in exact proportions to the configured transmit-rate

percentage. Therefore, when a physical interface operates in a PIR-only mode, this might cause imprecise scheduling results. [PR1200013](#)

- On MX Series platforms, the mspmand process might crash on the MS-MPC with XLP B2 chip (e.g.REV17). The exact trigger is unknown. It is usually seen with 70% to 90+% CPU load conditions. [PR1200149](#)
- GUMEM errors for the same address may continually be logged if a parity error occurs in a locked location in GUMEM. These messages should not be impacting. The Parity error in the locked location can be cleared by rebooting the FPC. [PR1200503](#)
- MS-MPC/MS-MIC: MSPMAND generates core files when an encrypted packet is received out of the range of replay-window size. The issue might occur in peak loads where by encrypted packets received, out of order due to drops in the network. [PR1200739](#)
- Dynamic firewall filter programs incorrect match prefix on the Packet Forwarding Engine [PR1204291](#)
- Packet Forwarding Engine may install next-hop incorrectly and cause traffic loss, if there is a next-hop policy pointing to a IPv6 address which need to be resolved. [PR1204653](#)
- If send upstream and downstream IPv4+IPv6 traffic for PPPoE subscribers, mirrored traffic loss would be seen. [PR1204804](#)
- VC link "last flapped" timestamp is reset to "Never" on the new backup Routing Engine after MX VC global GRES switchover. [PR1208294](#)
- The cpdd daemon might core and restart on the subscriber scenario with CPCD (captive-portal-content-delivery) service configured. [PR1208577](#)
- On MX Series running Tomcat release, if route-suppression is configured for access/access-internal routes as well as destination L2 address suppression is configured for the subscriber, bogus destination MAC would be generated for the subscriber. [PR1209430](#)
- The logic to calculate the IPsec phase2 soft lifetime has been changed in 14.2R6, resulting in an interop issue in certain scenarios. A hidden configuration statement is provided as part of this PR which will revert the soft lifetime logic to the one used in 11.4 release. [PR1209883](#)
- BGP PIC Installs multiple MPLS LSP next hops as Active instead of Standby in Packet Forwarding Engine. This can cause a routing loop. [PR1209907](#)
- On MX series routers, when configuring the dynamic access routes for subscribers based on the Framed-Route RADIUS attribute, the route will be created on the device, however, the it will be installed as an access-internal route instead of access route if it has /32 mask length. [PR1211281](#)
- Inline Jflow - Sequence number in flow data template is always set to zero on MPC5E and above line card type [PR1211520](#)
- On T-series platforms, if interfaces from FPC Type 4 and FPC TYPE 5 are configured together in one VPLS routing instance, incorrect TTL might be seen when packets go through the VPLS domain, for example, packets received via one FPC TYPE 4 might be forwarded to other FPC type 4 with incorrect TTL. The incorrect TTL could cause

serious VRRP issue. When VRRP is enabled, after one CE sends the VRRP advertise packets with TTL value 255, other CE might receive the VRRP packet with TTL value 0 and therefore discard these VRRP packets. As a result, the VRRP status in both CE becomes Master/Master. [PR1212796](#)

- The MS-MPC/MS-MIC service cards can encounter a core when using certain ALGs or the EIM/EIF feature due to a bad mapping in memory. [PR1213161](#)
- When FPC Type 5 - 3D cards run into over-temperature condition, in T4000 router. It is possible that under certain circumstances: - chassisd will declare the over-temperature condition and by default the router will shut down in 240 minutes. - Over-temperature SNMP trap (jnxOverTemp) are not sent to external NMS. [PR1213591](#)
- MX-VC: All VCP interface experiences tail-dropped as result of configuration conflict. It is a good idea to reference documentation and customize the COS associated with VCP interfaces. In this scenario customer has configured a corresponding xe-n/n/n interface with just a description to denote that port is dedicated to VCP. Problem is the resource calculation is impacted and reports smaller queue-depth maximum values when both network interface xe-n/n/n and vcp-n/n/n are defined. Issue is more likely to occur with dynamic modification add/delete of vcp interfaces with a corresponding network interface xe-n/n/n configured. > show interfaces queue vcp-5/3/0 | match max Maximum : 32768 Maximum : 32768 Maximum : 32768 Maximum : 32768
[PR1215108](#)
- If zero length interface name comes in the SDB database, on detection of a zero length memory allocation in the SDB database, a forced rpd crash would be seen. [PR1215438](#)
- On Junos OS Release 15.1R3 and later MX Series platform release, if DHCPv4 or DHCPv6 subscriber is configured and the subscriber joins more than 29 multicast groups, the line card might crash. [PR1215729](#)
- Incorrect source MAC used for PPPoE after underlying AE is changed [PR1215870](#)
- Prior to this fix for LI releases, parameterized family inet filter with term matching on address with non-contiguous mask will result in CLI syntax error which would fail subscriber login or CoA requests. [PR1215909](#)
- The AMS interface is configured in warm-standby mode when fail-over occurs a percentage of the traffic might fail to get NAT. The issue is after the failover the internal mappings driving traffic back to the service PIC might fail. [PR1216030](#)
- If RS/RA messages were received through an ICL-enabled(MC-AE) IFL, packet loss would be seen and last for a while. [PR1219569](#)
- The bbe-smgd core occurred in bbe_autoconf_if_l2_input when DHCP client generates ARP. [PR1220193](#)
- During CoA request there are no changes on schedulers. Requests are received successfully, but no changes from CoS side. [PR1222553](#)
- Due to a defect related to auto-negotiation in a Packet Forwarding Engine driver, making any configuration change to interface in MIC "3D 20x 1GE(LAN)-E,SFP" might lead to interface flapping. [PR1222658](#)

High Availability (HA) and Resiliency

- In PPP environment with access-internal and multiple routing instances, after restart rpd process, the access-internal route might disappear. [PR1174171](#)
- Backup routing engine might restart unexpectedly due to memory leak after switchover. [PR1198005](#)

Infrastructure

- With 13.3 releases using Ericsson/ Juniper EPG platforms, some session PIC C-PIC cards might experience some race condition resulting into kernel vmcores, following by reboot (failover to spare C-PICs) due to soft-update BSD enabled in some partitions of the Routing-Engine. The Softdeps on freebsd is not used any longer in freebsd6 where the fix includes disabling it on all Junos OS partitions. [PR1174607](#)
- From Junos OS Release 15.1 and later, smart error message of Unigen SSD may be seen. Smartd reads SSD attributes and checks on 197-current-uncorrectable, 198-offline-uncorrectable by default. To Unigen, 198 is not = Offline-Uncorrectable, it is 'Total Count of Read Sectors'. As it is Total-Read, such attribute(198) always carries value and smartd reports it as 'Offline Uncorrectable Error'. [PR1187389](#)
- The statistics info of em0 is 0 when checking by SNMP or CLI show command. [PR1188103](#)

Interfaces and Chassis

- In a VPLS scenario the flood NH for the default mesh group might not be programmed properly. A complete black-holing for the VPLS instance would be seen as a consequence. [PR1166960](#)
- In previous release, only IEEE classification is supported for CFM OAM packets. In the fix, we will support 802.1AD based filter for CFM OAM packets. when Linktrace and loopback requests are received in MX, 802.1p bits is used to determine the forwarding class and queue for response or linktrace request forwarded to next router, this cause these PDUs are put to wrong queue when input-vlan-map pop is present because received PDU doesn't carry 802.1p bits. In the fix, we will use incoming forwarding class to determine the 802.1p priority and outgoing forwarding class and queue for new generated response or link trace requests. [PR1175951](#)
- On dual Routing Engine system, if master Routing Engine is running Junos OS 13.3R9/14.1R7/14.2R5/15.1R3/15.2IB or later, backup Routing Engine is running Junos OS prior to 13.3R9/14.1R7/14.2R5/15.1R3/15.2IB, a major alarm is raised. This is cosmetic and can be safely ignored. Please upgrade backup Routing Engine to the same release with master Routing Engine to avoid the issue.

[illegible]

- In the hsl2 toolkit, there is a process which periodically checks the ASICs which communicate through it. Due to a bug in the toolkit code, the process used devalidate the very ASIC that it used to process, due to which the crash happens. [PR1180010](#)

- When there is a configuration change about OAM CFM, cfmd memory leak is observed and sometime also might trigger cfmd crash info as follows. Following messages are observed: /kernel: Process (44128,cfmd) has exceeded 85% of RLIMIT_DATA: used 378212 KB Max 393216 KB [PR1186694](#)
- The jpppd might crash with a core dump due to memory heap violation associated with processing MLPPP requests [PR1187558](#)
- If "filter" configuration statement is present in PPPoE traceoptions configuration, the resulting log file will contain only part of messages about establishment of the interesting PPPoE session, but will contain information related to other sessions established at the moment [PR1187845](#)
- SLR's/DMR's are not getting classified to Forwarding Class when CCM configured on AE with member links from NG MPC card. [PR1189254](#)
- In OAM CFM (connectivity-fault-management) scenario on AE interfaces with maintenance-domain level (for example: 3) configuration, when sending OAM CFM LBM messages with level which is smaller than configured level to ingress interface of VPWS with QinQ encapsulation, they are not dropped by ingress PE. [PR1191818](#)
- MAC addresses are incorrectly assigned to interfaces by the MX-VC SCC (global) chassisd daemon, leading to duplicate addresses for adjacent FPCs. [PR1202022](#)
- A CFMD core will be generated upon commit if the following conditions are met: * CFM is configured * On mis-configuration of icc format for MA (e.g. ICC name-format does not start with a character) [PR1202464](#)
- For the duration of GRES, if an async message for RTTABLE is received at DCD during initialization, it might result in unexpected state changes, the traffic forwarding might be affected. This is a timing issue, it is hard to reproduce. [PR1203887](#)
- When configuring "vlan-tags" for any interface, if the interface configuration is changed continually, the dcd process might memory leak. If the memory is exhausted, the dcd process might crash. [PR1207233](#)
- When VRRP is configured on IRB interface with scaling configuration (300k lines), in corner case, handles might not be released appropriately after their use is over. As a result of that, memory leak on vrrpd might be seen after configuration commit. [PR1208038](#)
- Access-internal route not installed for Dual Stack subscriber terminated in VRF at LNS with on-demand-ip-address [PR1214337](#)
- During L2TP session establishment on MX LAC, if CPE attempts to negotiate MRU higher than 1492 bytes, spurious MRU of 1492 bytes is included into the Last Received ConfReq AVP in ICCN packet. [PR1215062](#)
- In ppp subscriber scenario, if jpppd process receives the reply message from radius/tacplus server which has character of %, it might cause jpppd to crash. [PR1216169](#)

Layer 2 Ethernet Services

- In DHCP environment, if interface is deleted and recreated in single commit, the duplicate DHCP subscriber is not getting bound. [PR1188026](#)

- If a client sends a DHCP Request packet, and Option 55 includes PAD option (0), a DHCP ACK will not be sent back to the client. [PR1201413](#)
- In dhcp relay environment, when delay-authentication and proxy mode are configured at same time. Jdhcpcd may core due to NULL session ID. [PR1219958](#)

Multiprotocol Label Switching (MPLS)

- In the following scenario where 1) The PHOP link goes down and the router becomes MP for a LSP. 2) After some time, NHOP link for the same LSP goes down. The router becomes PLR for the same LSP. So effectively, the router is both MP and PLR for the same LSP. In this scenario, the router sends incorrect PathErr message for the backup MP PSB. It sends "Bad strict route" PathErr instead of "Tunnel local repaired" PathErr. [PR1132641](#)
- Due to Junos OS Release 15.1 enabling process rpcbind in FreeBSD by default, port 646 might be grabbed by rpcbind on startup, which causes LDP sessions failing to come up. [PR1167786](#)
- RSVP signalled p2mp sub-LSP with atleast 1 or more sub-LSPs in a down state might not get re-optimized in the event of a transit core link going down. If there are no sub-LSPs in a down state at the time of re-optimization then this issue won't be seen. This can cause traffic drop over the sub-LSP which are carrying traffic which are unable to get re-optimized. This PR addresses this issue. [PR1174679](#)
- On Juniper devices with "link-protection" configured and with/without "optimize-adaptive-teardown p2p" configured, rpd might crash after link flap. [PR1186003](#)
- With a high degree of aggregation and a large number of next hops for the same route, ldp may spend too much CPU updating routes due to topology changes. This may result in scheduler slip and ldp session timing out. [PR1192950](#)
- Packets will be out-of-order if they are Router Engine(RE) generated and go over unilist/ECMP. [PR1193697](#)
- Changing the configuration under both [protocols pcep] and [protocols mpls lsp-external-controller] might trigger rpd to crash due to a race condition. [PR1194068](#)
- If LDP neighbor relationship is over unnumbered interface, then flapping interface, the LDP will fail to advertise label binding. [PR1202071](#)
- With two Routing Engines and ldp export policy or l2-smart-policy configured. rpd on the backup Routing Engine may crash when ldp is trying to delete a filtered label binding. [PR1211194](#)

Network Management and Monitoring

- A trailing newline was erroneously added to the `$$message` variable, this had undesirable effects for some use cases when using the 'event-options policy `<>`' then `execute-commands` stanza. The fix escapes any newline chars which mitigates the issue. [PR1200820](#)

Platform and Infrastructure

- If IGMP snooping is enabled in a routing-instance (RI), in a very rare condition, the IGMP packets received in this RI might get dropped by firewall filter configured on loopback interface in master instance, which leads to multicast blackholing. [PR1092494](#)
- Preventing an issue where one could end up with two `<Junos: comment>` entries under the `[interfaces]` stanza. [PR1102086](#)
- In software versions which contain PR 1136360's code changes on MX-VC systems, when J-Flow is not configured and equal-cost multipath (ECMP) load-balanced routes occur, the linecards may stop forwarding packets after logging any of the below errors prior to possible linecard restart or offline:

- PPE Thread Timeout Trap. - PPE Sync XTXN Err Trap. - Uninitialized EDMEM Read Error. - LUCHIP FATAL ERROR. - `pio_read_u64()` failed.

(A possible workaround is to configure J-Flow and restart all linecards.)

In software versions which do not contain PR 1136360 solution, on MX Series Virtual Chassis (MX-VC) with "virtual-chassis locality-bias" configured, when equal-cost multipath (ECMP) load-balancing is occurring in the VC system, multicast streams and flooded Layer 2 streams may be duplicated or lost. Disabling "virtual-chassis locality-bias" from the configuration will eliminate the problem. [PR1104096](#)

- Configuring one group with configuration of routing-instances and applying this group under routing-instances, then the `rpd` process will crash after executing "deactivating/activating routing-instances" commands. As a workaround, you can avoid using "apply-groups" under routing-instances hierarchy. [PR1109924](#)
- On MX Series platform with MPC6E linecard, MPC6 only has 2 PICs (PIC number 0/1), if we try to configure an `si` interface with PIC number beyond range (PIC number 2) on MPC6E, it might crash, and traffic forwarding might be affected. [PR1160367](#)
- In CoS environment with shaping-rate configuration under interface, if flapping that CoS interface, the shaping-rate function does not take effect. As a workaround, please deactivate/activate interfaces to avoid the issue. [PR1163147](#)
- Because of an internal timer referring Time in Unix epoch (UNIX epoch January 1, 1970 00:00:00 UTC) value getting wrapped around for every 49 days, flows might get stuck for more than the period of active/inactive time out period. The number of flows that get stuck and how long they get stuck can not be deterministic exactly, which depends on the number of flows at the time of timer wrapping around. [PR1173710](#)
- This PR fixes an FD (file descriptor) leak problem in MGD process when `netconf traceoptions` are set. If `<commit> rpc` is executed via `netconf` session, there is an FD leak in the corresponding MGD pid. [PR1174696](#)

- On MX2020/2010, chassisd file rotation on commit check will cause the trace file to be stuck and no other operational chassisd events will be logged until chassisd restart. [PR1177625](#)
- When graceful Routing Engine switchover (GRES) is configured, the ksyncd crashes on backup Routing Engine (RE) if a VPN static route has a network address as a next-hop. This causes that the backup Routing Engine is not ready for graceful switchover. [PR1179192](#)
- The issue happens after GRES. If commit on the new master during the config sync from the old master, commit might fail. [PR1179324](#)
- In IPv6 sampling environment, if flapping IPv6 routes frequently, in rare condition, due to a software defect, free of route node is not deleting it from radix node, so the Packet Forwarding Engine might crash. This is a corner case, it is hard to reproduce. [PR1179776](#)
- On MX platform with LU chipset such as MPC1/ MPC2/ MPC3E/ MPC4E/MPC 3D 16x10GE or T platform with FPC type 5, if one interface is applied COS schedulers with transmit-rate percent and rate-limit parameter, then for pseudowire traffic, the traffic transmit-rate percent is not correct. [PR1180427](#)
- If igmp snooping is configured in a VPLS routing instance and the VPLS instance has no active physical interfaces, multicast traffic arriving from the core might be send to the Routing Engine. As a result, host queues might get congested and it might cause protocol instability. As a workaround, configure a dummy activate interface in the VPLS routing instance can avoid this issue. [PR1183382](#)
- On MX2K, the 'commit full' operation, or committing configuration under 'system' stanza (such as root-authentication and fxp0 interfaces) can cause transient Fan check Major alarm and Fan full speed. The Fan Tray spins at full speed for a while, then goes back to normal with clearing the alarm. The Fan check alarm and corresponding snmp trap are temporal, and they can be safely ignored.

```

user@MX2K> show chassis alarms 2 alarms currently active Alarm time Class
Description 2016-05-17 19:49:57 JST Major Fan Tray X Failure 2016-05-17 19:49:57 JST
Major Fan Tray Y Failure
usr@MX2K> show chassis environment Class Item Status
Measurement Fans Fan Tray X Fan 1 Check Fan Tray X Fan 2 Check Fan Tray X Fan 3
Check Fan Tray X Fan 4 Check Fan Tray X Fan 5 Check Fan Tray X Fan 6 Check Fan
Tray Y Fan 1 Check Fan Tray Y Fan 2 Check Fan Tray Y Fan 3 Check Fan Tray Y Fan 4
Check Fan Tray Y Fan 5 Check Fan Tray Y Fan 6 Check

```

When MPC9E is installed in MX2K, the Fans usually keep around 6K rpm, and the fan speed control is frequently done by the Junos OS software. In this situation, when all daemons are re-evaluated (by commit full or config change under system stanza), the software bug causes the fan status to be checked within quite small period, then the Junos OS software recognizes that the fan is faulty because the fan speed has not reached the target speed yet when the fan status is checked within the small period. After the fan alarm is detected, the fans are expected to start working with full speed to cool the system components.

The fan status check logic is fixed by this PR. The fan status is checked after the fan speed is stabilized, hence we do not see this transient fan alarm. [PR1185304](#)

- In a very rare scenario, during TAC accounting configuration change, auditd daemon crashes due to a race condition between auditd and its sigalarm handler. [PR1191527](#)
- VPLS: FPC CPU goes high for several minutes when mac/arp are learnt via lsi interfaces. The FPC CPU goes high during the learning phase and issue can be seen with various triggers that result in mac/arp re-learning e.g. mac flush, FPC reboot or link flap resulting in mac flush etc. For agent smith cards (MPC 3D 16x 10GE), the CPU may remain high for upto 30 minutes on learning/re-learning of 10k arp/mac via lsi interfaces Problem is only seen if there are ARPs learnt in bulk over lsi interfaces. [PR1192338](#)
- Insertion of an offlined MPC6E into the MX2K chassis can cause the FPC Temp sensor to detect transient "WARM TEMP" condition, and the chassis FAN in the same zone goes to high speed.

*** messages ***

```
Jul 12 18:10:17.698 MX2K-re0 chassisd[xxxx]: CHASSISD_SNMP_TRAP7: SNMP trap
generated: FRU insertion (jnxFruContentsIndex 7, jnxFruL1Index 3, jnxFruL2Index 0,
jnxFruL3Index 0, jnxFruName FPC: MPC6E 3D @ 2/*/*, jnxFruType 3, jnxFruSlot 2)
MX2K-re0> show chassis zones |refresh 2 ---(refreshed at 2016-07-12 18:10:18 JST)---
ZONE 0 Status Driving FRU FPC 2 Temperature 63 degrees C / 145 degrees F Condition
WARM TEMP <----- Warm temp is detected Num Fans Missing 0
Num Fans Failed 0 Fan Duty Cycle 27 ZONE 1 Status Driving FRU SFB 5 SFB-XF2-Zone1
Temperature 59 degrees C / 138 degrees F Condition OK Num Fans Missing 0 Num
Fans Failed 0 Fan Duty Cycle 27 ---(refreshed at 2016-07-12 18:10:20 JST)--- ZONE 0
Status Driving FRU FPC 2 Temperature 63 degrees C / 145 degrees F Condition WARM
TEMP <----- Warm temp is detected Num Fans Missing 0 Num Fans
Failed 0 Fan Duty Cycle 27 ZONE 1 Status Driving FRU SFB 5 SFB-XF2-Zone1
Temperature 59 degrees C / 138 degrees F Condition OK Num Fans Missing 0 Num
Fans Failed 0 Fan Duty Cycle 27 ---(refreshed at 2016-07-12 18:10:22 JST)--- ZONE 0
Status Driving FRU FPC 2 Temperature 63 degrees C / 145 degrees F Condition OK
Num Fans Missing 0 Num Fans Failed 0 Fan Duty Cycle 27 ZONE 1 Status Driving FRU
SFB 5 SFB-XF2-Zone1 Temperature 59 degrees C / 138 degrees F Condition OK Num
Fans Missing 0 Num Fans Failed 0 Fan Duty Cycle 27
```

```
Jul 12 18:10:27.489 MX2K-re0 chassisd[xxxx]: Fan Tray 0: zone 0 fan_speed current
27% target 50% raising ratio 0.80 (linear) FPC 2 temp 72 last 72 WTC 55 WT 60 high
limit 75 i2c_ratio 0.80 Jul 12 18:10:27.490 MX2K-re0 chassisd[xxxx]: Fan Tray 0: set
fan_speed to 50% cfg_speed 50% (linear) Jul 12 18:10:27.492 MX2K-re0 chassisd[xxxx]:
Fan Tray 1: zone 0 fan_speed current 27% target 50% raising ratio 0.80 (linear) FPC
2 temp 72 last 72 WTC 55 WT 60 high limit 75 i2c_ratio 0.80 Jul 12 18:10:27.492
MX2K-re0 chassisd[xxxx]: Fan Tray 1: set fan_speed to 50% cfg_speed 50% (linear)
Jul 12 18:10:47.517 MX2K-re0 chassisd[xxxx]: Fan Tray 0: zone 0 fan_speed current
50% target 27% falling ratio 0.00 (linear) SFB 2 SFB-XF0-Zone0 temp 63 last 63
WTC 70 WT 75 high limit 90 i2c_ratio -0.60 Jul 12 18:10:47.517 MX2K-re0 chassisd[xxxx]:
Fan Tray 0: set fan_speed to 27% cfg_speed 27% (linear) Jul 12 18:10:47.519 MX2K-re0
chassisd[xxxx]: Fan Tray 1: zone 0 fan_speed current 50% target 27% falling ratio 0.00
(linear) SFB 2 SFB-XF0-Zone0 temp 63 last 63 WTC 70 WT 75 high limit 90 i2c_ratio
-0.60 Jul 12 18:10:47.520 MX2K-re0 chassisd[xxxx]: Fan Tray 1: set fan_speed to 27%
cfg_speed 27% (linear) PR1193273
```

- A rare VMCORE can occur caused due to process limit being breached by too many RSHD children processes being created [PR1193792](#)
- After system boot up or after PSM reset we may see "PSM INP1 circuit Failure" error message [PR1203005](#)
- When a Netconf <get route information> RPC is executed for all routes via ssh transport session and the session is terminated before all the route information is retrieved, the MGD process and RPD daemon will cause high CPU utilization for an extended period of time. Example of issues caused by this high CPU utilization for an extended period is as follow: BGP neighbors holddown timer expires and become ACTIVE OSPF adjacencies reset during database exchange OSPF LSA retransmissions events on neighboring nodes due to missing ACKs LDP sessions time out non distributed BFD sessions being reset due to missing keepalives [PR1203612](#)
- If Inline JFlow is configured in scaled scenarios, Inline JFlow Sampler route database is taking huge time to converge. [PR1206061](#)
- When "commit confirmed" is used after performing some changes, and an empty commit is performed to confirm the changes, the previous changes related processes will be notified again which is unnecessary. It might cause session/protocol flap. [PR1208230](#)
- If a Unicast or Multicast source sends a fragmented packet (a packet which exceeds the MTU of its outgoing interface) to the router and it needs to resolve the destination route, then only the first fragment of the packet is sent when the route is resolved. [PR1212191](#)
- On MX2K, MIC output is seen when there is no MIC in MPC under "show chassis hardware detail".

Steps to reproduce the issue: 1. offline MPC 2. physically remove MPC 3. physically remove MIC from the MPC 4. reinsert MPC 5. online MPC

```
usr@MX2K> show chassis hardware detail |find fpc FPC 0 REV 68 750-044130
ABDA1879 MPC6E 3D CPU REV 12 711-045719 ABDA1735 RMPC PMB MIC 0 REV 14
750-049457 ABCY5322 2X100GE CFP2 OTN >>>>>>> No MIC inside MIC 1 REV 26
750-046532 ABCZ3853 24X10GE SFPP >>>>>>>>No MIC inside XLM 0 REV 13
711-046638 ABDA1859 MPC6E XL XLM 1 REV 13 711-046638 ABDA1787 MPC6E XL
PR1216413
```

- This rmopd core was caused by the NULL pointer in SW function. [PR1217140](#)
- When any MPC line card is offlined, it goes offline via all offline flows and connection is cleaned, but in the end of the offline flow, somehow it delays powering off the line card. The chasd powers off the MPC via I2cs write the respective power registers, but in hardware it is not really powering off. As a consequence, since MPC is still power-on but connection is down, it will try to reconnect, then start to come up automatically within 10 secs. It occurs sometimes (not all the times). [PR1222071](#)

Routing Policy and Firewall Filters

- With rib-groups configured for importing routing information to multiple routing tables, unexpected route refresh might happen when committing configuration change, due to a defect in code related to secondary table list handling. [PR1201644](#)

- From Junos OS Release 15.1, memory leak on policy_object might be observed if the configuration of policies is added and deleted in high frequency. Not all policies make memory leak, and only the container policy referred in policy statement hits this issue: the "from" in policy invokes the terms which is defined in policy-options, e.g. community, as-path, prefix-list. This is the configuration example. set policy-options prefix-list pl set policy-options policy-statement from prefix-list pl [PR1202297](#)

Routing Protocols

- Junos OS exhibits two different next-hop advertisement behaviors for MP_REACH_NLRI on a multi-hop eBGP session, based on whether it is loopback peering or physical interface peering. When the routers are peering on their loopback, only the global IP of the interface (lo0) is advertised, whereas when the routers are peering through the physical interface, both global and link-local address are advertised as the NHs. [PR1115097](#)
- When BGP speaker has multiple peers configured in a BGP group and when it receives the route from a peer and re-advertises route to another peer within the same group, MIB object "jnxBgpM2PrefixOutPrefixes" to the peers in the same group reports the total number of advertised prefixes in the group. MIB value "jnxBgpM2PrefixOutPrefixes" is defined as per peer basis but it looks as if it is per group basis. As a workaround, we can get the number of advertised prefixes from CLI command "show bgp neighbor" instead. [PR1116382](#)
- When Bidirectional Forwarding Detection (BFD) is configured, after changing the MTU (between 1514 and 9192) of physical interface (IFD) where the BFD session is located, 2 issues might be seen as below. Issue 1: after link flapping, the BFD session may not come up due to incorrect mapping. Issue 2: there might be stale BFD sessions. This issue may also be seen when changing the interval from aggressive to a very less aggressive interval (e.g. change to 2 sec). [PR1116666](#)
- On Junos OS based products, changes in routing-instance, like changing route-distinguisher or routing-option changes in some corner cases might lead to rpd crash. As a workaround always deactivate routing-instance part that is to be changed before committing the changes. [PR1134511](#)
- When we have a route received from different eBGP neighbors, for this specific route, if all BGP selection criteria is matching, we will end up using router ID. As this is eBGP route, so BGP will use active route as the preferred one. Now if this specific route flapped with sequence from the non-preferred to the preferred path, RPD will run the path selection. During RPD path selection we might generate a core file. This issue has no operational impact, also a workaround is available to avoid this issue. [PR1180307](#)
- Please refer to the following topology. If the opposite Router's interface "A" is down by "disable/deactivate/delete" configuration, BFD timeout detection might be long delay. Topology +-----+ | DUT | OSPF | |-----+ +-----+ | A | | | | | +-----+ OSPF(p2p) | | R2 | bfd | | | | +-----+ | | V intf A | | +-----+ | | R1 |-----+ | | OSPF +-----+ [PR1183353](#)
- If we have post-policy BMP configured & import policy rejects the route making it hidden, we will still periodically send this Unreachable Prefix to the BMP station.

May 17 15:45:05.047931 bmp_send_rm_msg called, found post-policy prefix 101.66.66.66/32, peer 10.0.1.1 (External AS 65101), station BMP_STATION_2. May 17 15:45:05.047943 import policy rejected post-policy prefix 101.66.66.66/32, peer 10.0.1.1 (External AS 65101), station BMP_STATION_2. May 17 15:45:05.047986 generating post-policy delete for prefix 101.66.66.66/32, peer 10.0.1.1 (External AS 65101), station BMP_STATION_2. May 17 15:45:05.048001 BMP: type 0 (RM), len 76, ver 3, post-policy, for Peer 10.0.1.1, station BMP_STATION_2. May 17 15:45:05.048018 Peer AS: 65101 Peer BGP Id: 10.0.1.1 Time: 1463492684:0 (May 17 13:44:44) May 17 15:45:05.048027 Update: message type 2 (Update) length 28. May 17 15:45:05.048034 Update: Unreachable prefix data length 5. May 17 15:45:05.048047 Update: 101.66.66.66/32 [PR1184344](#)

- Any configuration change can cause deletion of a firewall filter created for a routing instance if the flowspec routes in that instance are imported using rib-group, and there is no "inet-vpn flow" address family configured and the routing instance does not have any BGP group configured with "inet flow" address family. [PR1185954](#)
- On the RSVP LSP scenario with ISIS TE configured, memory leak might happen in rpd and Packet Forwarding Engine after the LSP re-optimization, and this might cause FPC crash. [PR1187395](#)
- The rpd might crash when printing the socket address of type inet6 flow address family while the buffer is not sufficient to print decimal number. [PR1188502](#)
- Multicast routing table displays inconsistent MoFRR state after activating/deactivating MoFRR. This is a cosmetic issue and has no impact on traffic. [PR1194729](#)
- On executing "show task replication" command, IS-IS could be shown as "Complete" if IS-IS is not configured on the device. If IS-IS is configured, the replication will be shown correctly (NotStarted/InProgress/Complete). No other functionality impacted. [PR1199596](#)
- The VRF related routes which are leaked to the global inet.0 table and advertised by the access routers are not being advertised to global inet.0 table on the core. [PR1200883](#)
- With nonstop-routing (NSR) enabled, all running protocols include PIM and NG-MVPN will be replicated, if NSR is disabled only under PIM "set protocol pim nonstop-routing disabled", this will remove both PIM and NG-MVPN from replicated list, then adding PIM NSR again by "delete protocol pim nonstop-routing disabled" will not work as expected and PIM will not be added. [PR1203943](#)
- In a situation which a BGP route is resolved using a secondary OSPF route which is exported from one routing-instance to another routing-instance. If the BGP route is being withdrawn while the OSPF route is deleted, rpd might restart unexpectedly. [PR1206640](#)
- BGP routes are rejected as cluster ID loop prevention check fails due to a mis-configuration. But when the mis-configuration is removed BGP routes are not refreshed. The fix of this issue will send a soft route refresh dynamically when a cluster ID is deleted. [PR1211065](#)

- If a NSR enabled router is providing graceful restart support for a restarting peer, and the standby is unconfigured, then rpd may core on the standby during the shutdown. [PR1212683](#)
- EBGp peer may remain "Idle" at NSR backup-Routing Engine, after Interface-down event [PR1215855](#)

Services Applications

- On MX platform, when using MS-MPC, the "idpd_err.date" error message is filling var/log. Please refer to KB30743 for details. [PR1151945](#)
- During "commit synchronize" operation, when commit gets executed on backup Routing Engine, system is idling for 10 seconds after the following operation (can be observed with "commit synchronize | display detail"): 2016-07-07 10:30:04 CEST: Spawning IPsec Key Management daemon to check new configuration This slows down the whole commit process exactly by 10 seconds. Issue can only be seen when IPsec is configured and, therefore, IPsec Key Management daemon (kmd) is running (needed by configuration). [PR1185504](#)
- When using MS-DPC under heavy load condition (e.g. with about 7m flows) with deterministic NAT and port block allocation (PBA) scenario, in rare condition, MS-DPC crash may occur due to memory issue. [PR1186391](#)
- Attempting to ping a subscriber address from the L2TP LNS CLI will fail. [PR1187449](#)
- Issue happens in specific corner cases and Acceptable workaround is available. If we bring down the complete subscriber and bring it back up again. Family bring up will work. [PR1190939](#)
- When using NAT on the MX Series the FTP ALG fails to translate the PORT command when the FTP client using Active Mode requests AUTH(SSL-TLS) and the FTP server does not use AUTH [PR1194510](#)
- When MS-PIC is running on T640/T1600/T4000, the number of maximum service sets is wrongly limited to 4000, instead of 12000. This might impact in scaled service (IPsec, IDS, NAT, Stateful firewall filter, etc) environment. [PR1195088](#)
- After upgrading M series router (LNS) to 15.1R4.6, it was observed that L2TP sessions are not coming up due to PPP CHAP authentication failure. L2TP control messages are sent/received and tunnel id is obtained. PPP LCP is also successful. During PPP CHAP phase only Challenge and Response messages are present and then L2TP CDN is initiated. [PR1201733](#)
- When configuring Network Address Translation (NAT) service, the service route is still available in route table even after disabling service interface. Any types of service interfaces (except ams- interface) that supports NAT might be affected. [PR1203147](#)
- On MX series with L2TP configured, for some reason the L2TP packet in ICRQ retransmission message is set to incorrect value, and this causes frequent L2TP session flaps. [PR1206542](#)
- On MX Series routers with subscriber management feature enabled used as a LAC (L2TP Access Concentrator), a small amount of memory leak is leaked by jl2tpd process on the backup Routing Engine when subscriber sessions are logged out. [PR1208111](#)

Subscriber Access Management

- In DHCP relay scenario, DHCP relay binding might get stuck in "RELEASE(RELAY_STATE_WAIT_AUTH_REQ_RELEASE" state due to the LOGOUT Request is not processed correctly by authentication manager process (authd) if there were multiple attempts to activate Lawful Intercept (LI) for this DHCP subscriber using RADIUS change of authorization (CoA) packets in quick succession. [PR1179199](#)
- If aborting "test aaa ppp" command with Ctrl-C, due to a software defect, when subscriber logout, the system does not wait for logout response, subscriber is immediately removed. Because of this, dfwd daemon is not able to clear filters in time and results in stale entries. The stale info might affect subscriber login and logout. [PR1180352](#)
- In the event, such as JSRC re-sending a PPR with a policy-install for an already installed policy or policy-remove for a non-existing policy (resulting if the SRC goes down after issue the PPR but before receiving or preserving the response), the outcome of the processing is to "do-nothing" which results in a different code path. [PR1189020](#)
- On EX2200/EX3300 series switches configured dhcp-local-server, it brings up a few (say 6 or more) or all interfaces which is under dhcp-local-server hierarchy at once then the authd process continually core dumps causing the switch get in stuck and resulting in packet drop. [PR1191446](#)
- When destination-override is used (root@user# set system tracing destination-override syslog host <host ip>), the userAccess events are not sent to the external syslog server. [PR1192160](#)
- On MX series platform, when using RADIUS dynamic requests for subscriber access management, if the device detects that the CoA-Request it received is same with the one in processing progress, the router would send CoA-NAK packet back to the RADIUS server with incorrect code 122 (invalid request) wrongly, before sending CoA-ACK packet in response to the original CoA-Request that was being processed. In this case the router should ignore all RADIUS CoA-Request retries and respond only to the original CoA-Request packet. [PR1198691](#)
- Incorrect service-accounting name in radius accounting record if service activated by SRC [PR1206868](#)
- If RADIUS return Framed-route="0.0.0.0/0" to a subscriber terminated on Junos OS platform, this subscriber can not login due to authentication error. [PR1208637](#)
- On MX Series routers with subscriber management feature enabled, after GRES switchover "show network-access aaa statistics radius" CLI command display only zeros and "clear network-access aaa statistics radius" doesn't clear statistics as it should. It's a cosmetic issue and communication with Radius server is working fine, the only impact is that affected CLI commands do not work as expected. [PR1208735](#)
- If radius Primary-WINS (Juniper-ERX-VSA) is set as 0.0.0.0, subscribers is rejected by Authd and doesn't negotiate further. [PR1209789](#)
- Commit error: "Radius-Flow-Tap LSRI "" is in use by subscriber, cannot be removed from the configuration" might be seen after two consecutive GRES switchovers if a

subscriber with lawful intercept mirroring enabled was logged in before the switchovers. [PR1210943](#)

User Interface and Configuration

- When entering the "restart r" incomplete command in the CLI, the command "restart routing" is executed. It should throw an error like "error: invalid daemon: r". [PR1075746](#)
- Config database is locked by "root" user when trying to commit vpls circuit configs in "config exclusive" mode. [PR1208390](#)
- If user enter configuration mode with "configure exclusive" command, after configuration is automatic rollback due to commit un-confirmed, user still can make configuration changes with "replace pattern" command, the subsequent commit fails with "error: access has been revoked". After exit configuration mode, user fail to enter configuration mode using "configure exclusive" with "error: configuration database modified". [PR1210942](#)
- When persist-groups-inheritance is configured and you issue a rollback, it will be seen that the configuration is not propagated properly after a commit. [PR1214743](#)

VPNs

- In a multi-homed source topology in NG-MVPN (applicable to both inter-AS and intra-AS scenario), there are two problems: The first problem is Multicast (S, G) signaling doesn't follow RPF. When the routing table (mvprninstancename.inet0) has two routes, due to the policy configuration, the best route to the source is via the MPLS core, but Multicast (S, G) PIM join and NG-MVPN Type 7 both point to inactive route via local BGP peer. The second problem is when "clear pim join instance NG" is entered, the multicast forwarding entries are wiped out. [PR1099720](#)
- After a GRES with NSR enabled, in NG-MVPN scenario, on the new backup RE RPD is consuming more than 90% CPU. This issue happens rarely and it is not reproducible. [PR1189623](#)
- In BGP VPLS environment, sometimes we receive routes from BGP with invalid next-hop related information. In such scenarios, VPLS should treat them as bad routes and not send them to rpd infra for route resolution. Due to a software defect, the bad routes are passed to the route resolver, which might lead to rpd process crash. The routing protocols are impacted and traffic disruption will be seen due to loss of routing information. [PR1192963](#)
- With MVPN and NSR enabled, high CPU on backup Routing Engine might be seen. MVPN on backup Routing Engine is re-queuing c-mcast events for flows as it is unable to find phantom routes from master routing-engine. However as routes is not reaching from master Routing Engine so backup Routing Engine keeps trying causing high CPU triggered by rpd processing. [PR1200867](#)
- In MVPN mode SPT-only, the first multicast packet is lost when the source is directly connected to the PE. [PR1204425](#)

Resolved Issues: 15.1R4

- [Class of Service \(CoS\) on page 234](#)
- [Forwarding and Sampling on page 234](#)
- [General Routing on page 236](#)
- [High Availability \(HA\) and Resiliency on page 242](#)
- [Infrastructure on page 243](#)
- [Interfaces and Chassis on page 243](#)
- [Layer 2 Features on page 245](#)
- [MPLS on page 245](#)
- [Network Management and Monitoring on page 247](#)
- [Platform and Infrastructure on page 247](#)
- [Routing Policy and Firewall Filters on page 251](#)
- [Routing Protocols on page 251](#)
- [Services Applications on page 252](#)
- [Subscriber Access Management on page 253](#)
- [User Interface and Configuration on page 253](#)
- [VPNs on page 253](#)

Class of Service (CoS)

- When customers delete an IFL from an interface-set that has CoS applied to it and activate CoS profile directly on that IFL in one single commit, commit fails with an error. Commit goes through if they do it one by one, delete IFL from interface set, commit and then activate CoS on that IFL, commit. [PR1169272](#)

Forwarding and Sampling

- Configuration statement "interface-mac-limit" might be set to default value when activating "mac-table-size" on a VPLS routing instance. Restarting l2ald, reapplying the "interface-mac-limit" or changing to another value (set interface ge-3/1/0.0 interface-mac-limit 510) fixes the issue. user@router> show vpls statistics | match count Current MAC count: 0 (Limit 1024) << set to default value 1024 instead of the value set by interface-mac-limit [PR1025503](#)
- In some rare cases, SNMP might get Output bytes of Local statistics instead of the Traffic statistics when retrieving Output bytes of Traffic statistics on a logical interface. [PR1083246](#)
- When using MX Series-only features (gre decapsulate or payload protocol in IPv6), a change of policers or counters to an existing firewall filter using physical-interface-filter or interface-specific configuration statements will not be correctly detected by MIB2D. [PR1157043](#)
- Configuration container [protocols] [l2-learning] [global-mac-move] is made visible. The functionality under it are already supported but the command was hidden till now. [PR1160708](#)

- Configuration is restricted to include uid variables in variable expressions Please find the following example as below root@R1# show dynamic-profiles SERVICE-PROFILE variables input-filter { mandatory; uid-reference; } input-bw mandatory; output-filter { mandatory; uid-reference; } output-addr1 mandatory; output-addr2 mandatory; fin1-uid uid; fout1-uid uid; fout2-uid uid; policer1-uid uid; prefix1-uid uid; term-var equals "ifNotZero (\$output-addr1,'voice:###\$fout2-uid##':'###\$fout1-uid)"; root@R1# commit error: syntax error in profile SERVICE-PROFILE variable term-var error: syntax error in variables stanza in profile SERVICE-PROFILE error: foreign file propagation (ffp) failed. [PR1168994](#)

- This issue will be seen only when there are huge number of routes having different BGP NHs pointing to the same AS. Depending on the number of routes pointing to AS paths and also the difference in BGP NHs in the routes can shoot up the SRRD CPU consumption. In the real network this issue might not be seen often, as the number of AS paths will be huge and the routes referring these AS paths will be usually distributed among the AS paths. Even if the routes are pointing to the same AS, the impact would be lesser than the one seen in this PR. [PR1170656](#)

- When polling SNMP counters for MX series-Only firewall filters, MIB2D_RTSLIB_READ_FAILURE cosmetic error messages might get reported in syslog. [PR1173057](#)

- statistics-service daemon (pfed) experiences constant memory leak of 10 KB every 2 minutes when MobileNext package is installed: > show version Model: mx480 Junos: 14.1X55-D30.10 JUNOS Base OS boot [14.1X55-D30.10] <...> JUNOS MobileNext Routing Engine Software [14.1X55-D30.10] <<< this package. [PR1174193](#)

- Even if packets do not match firewall filter conditions, wildcard mask firewall filter might match any packets. << Sample config >>

```

----- set firewall family inet filter TEST-filter
term TEST1 from destination-address 0.0.0.255/0.0.0.255 <<<<< set firewall family
inet filter TEST-filter term TEST1 then count TEST1 set firewall family inet filter
TEST-filter term TEST1 then discard set firewall family inet filter TEST-filter term
TEST2 then accept ----- This is discard filter
for /24 prefix broadcast address. However it might discard other packets. PR1175782

```

- This is cosmetic issue. During sampling with jflow version 9, bfd packets from MPLS-TP were shown like as ip packets in "show services accounting aggregation template template-name XXX" command. (Actually, bfd packets info is not sampled by jflow.)
<< example >>

```

*****
lab@router-re0> show services accounting aggregation template template-name
mpls Src Dst Port/ Port/ Top MPLS MPLS MPLS Source Destination ICMP ICMP Label
Label 1 Label 2 Label 3 Address Address Type Code Proto TOS Address 299776 13 0
0.0.0.16 0.1.134.160 0 0 0 100.100.100.3 <<<<< bfd packet 299776 13 0 0.0.0.17
0.1.134.160 0 0 0 100.100.100.3 <<<<< bfd packet 299776 16 0 10.0.0.1 40.0.0.2 8
0 1 0 100.100.100.3 <<<<< ping 299792 16 0 40.0.0.2 10.0.0.1 0 0 1 0 100.100.100.1
<<<<< ping 299776 16 0 40.0.0.2 10.0.0.1 0 0 1 0 100.100.100.3 <<<<< ping
***** <
sample topology >>
*****

```

```
MPLS-TP(OAM, BFD) <-----> 10.0.0.1 40.0.0.2 sampling
[CE1]-----[PE1]-----[DUT]-----[PE2]-----[PE2] || [collector]
*****
PR1177876
```

- In Junos OS Release 15.1 and later, family vpls filter applied to ae-interface is not working. [PR1178743](#)
- SRRD daemon does not delete routes when the DELETE is received from RPD in few configuration cases. This results in build-up of memory in SRRD daemon and once SRRD reaches the limit, it crashes and restarts itself. This happens only when none of the SRRD clients (FPCs in Inline JFlow case and PICs in PIC based sampling) are interested in one or more families. Say, only IPv4 family is configured in all the clients and, IPv6 and MPLS families are not configured for Sampling in any of the clients. [PR1180158](#)

General Routing

- An EVPN with support for inter-subnet routing using an irb interface may experience a crash and restart of rpd, leaving a core file for analysis. In this case, EVPN MAC routes contain MAC+IP, and this IP/32 is installed in Routing Instance table on egress router. Core is triggered in the IP/32 route installation flow. There is no special trigger point-it is a timing issue with basic irb configurations. [PR992059](#)
- An inconsistency between JUNIPER-VPN-MIB and MPLS-L3VPN-STD-MIB with the number of interfaces for an routing-instance has been identified. For example with the following configuration: user@router-re0> show configuration routing-instances ri1 instance-type vrf; interface ge-2/0/8.10; interface lo0.10; route-distinguisher 65000:1; vrf-target target:65000:1; vrf-table-label; According to the MPLS-L3VPN-STD-MIB there are two interfaces in this routing-instance: MPLS-L3VPN-STD-MIB :: mplsL3VpnVrfAssociatedInterfaces: OID: 1.3.6.1.2.1.10.166.11.1.2.2.1.8 Description: Total number of interfaces connected to this VRF (independent of ifOperStatus type). {master} user@router-re0> show snmp mib walk 1.3.6.1.2.1.10.166.11.1.2.2.1.8 mplsL3VpnVrfAssociatedInterfaces.3.114.105.49 = 2 However according to JUNIPER-VPN-MIB there are three interfaces in this VRF: JUNIPER-VPN-MIB :: jnxVpnIfStatus OID: 1.3.6.1.4.1.2636.3.26.1.3.1.10 Description: Status of a monitored VPN interface. user@router-re0> show snmp mib walk 1.3.6.1.4.1.2636.3.26.1.3.1.10 jnxVpnIfStatus.2.3.114.105.49.733 = 5 jnxVpnIfStatus.2.3.114.105.49.754 = 5 jnxVpnIfStatus.2.3.114.105.49.774 = 5 The interfaces in the example are: {master} user@router-re0> show snmp mib walk 1.3.6.1.2.1.2.2.1.2 ifDescr.733 = ge-2/0/8.10 ifDescr.754 = lo0.10 ifDescr.774 = lsi.0 The fix for this issue adjusts this by removing the dynamic interface (in this case, lsi.0) from the interface list of JUNIPER-VPN-MIB. [PR1011763](#)
- The L2ald may crash after interface flap. [PR1015297](#)
- CoS scheduler names cannot be added or changed via service COA's. The schedulers can be added at subscriber login using client dynamic profiles. [PR1015616](#)
- When ps interface is configured using as anchor interface, a logical tunnel (lt) interface without explicit tunnel-bandwidth configuration (under 'chassis fpc <fpc number> pic <pic number> tunnel-services' configuration hierarchy), the ps interface is created

only in kernel, but not on Packet Forwarding Engine. In order to have ps interface in Packet Forwarding Engine, an explicit tunnel-bandwidth configuration is required. PR1042737 removes this restriction, and a ps interface may be anchored to an lt interface without explicit tunnel-bandwidth configured. [PR1042737](#)

- IPV6 RA is not including source link address option on ps.x pseudowire interfaces. [PR1049952](#)
- Wrong byte count was seen in the ipfix exported statistics packets for mpls flows. This issue is taken care now. [PR1067084](#)
- There are some configuration related functions in rpd and l2cpd that use special Memory API called Lite Pools. These pools when reset were not freeing control information related to the pool and hence resulting in a leak. This is not a day one issue. This bug was introduced in 15.1 when we reimplemented LIBTASK memory subsystem. This PR impacts all daemons using LIBTASK (including rpd) on all platforms provided memory lite pools are used by those daemons. [PR1071191](#)
- PCE-initiated LSPs are less preferred than locally configured LSPs. After this issue is fixed, PCE-initiated LSPs will have same preference as locally configured LSPs. [PR1075559](#)
- The Enhanced LAG feature is enabled in network-service enhanced-ip mode, but it is not supported in enhanced-Ethernet mode. [PR1087982](#)
- Certain VTY JNH commands (see description of this PR-1094955) on MX Series platforms will not decode properly, would need this PR fix. [PR1094955](#)
- On MX Series routers where MS-MIC or MS-MPC is inserted, certain combinations of fragmented packets might lead to an MS-MIC or MS-MPC coredump. [PR1102367](#)
- On MX Series platforms, in rare condition, if Packet Forwarding Engine sends wrong Packet Forwarding Engine id to chassisd as part of capability message, kernel might crash and some FPCs might be stuck in the present state, the traffic forwarding will be affected. This is a corner case, it is not reproduced consistently. [PR1108532](#)
- Fixed problem with "egress pfe unspecified" increase when bind dhcp relay (or fpc restart caused ospf connection lose. Not able to ping its neighbor, arp table is fine, got egress Packet Forwarding Engine unspecified). [PR1114132](#)
- ANCP is not supported in this release. Attempts to use ANCP related show commands will result in a timeout. [PR1121322](#)
- With IPv6 access route configured in dynamic profile, when the router receives IPv6 SOLICIT message which request only Prefix Delegation but no IPv6 address, the access route will not be installed successfully. [PR1126006](#)
- RPD crash might be seen during deletion of address family on an interface while rpf check is configured. [PR1127856](#)
- The rpd might crash when local-switching is configured with connection-protection for L2Circuit. This problem only occurs after setting composite-next-hop for L2Circuit using **set routing-options forwarding-table chained-composite-next-hop ingress l2ckt**. [PR1129940](#)

- When using Point-to-Point Tunneling Protocol (PPTP) Application Layer Gateways (ALG) on MS-MPC/MS-MIC, if running scaled number of PPTP sessions control and data sessions (e.g. 1M sessions) for long hours (e.g. more than 8 hours), when the traffic is stopped, the "Bytes used" field of the output of CLI command "show services service-sets summary" will show a randomly large value due to memory issue. [PR1131605](#)
- In a situation where both mirrored interface and mirrored destination are on MPC card and mirror destination interface is a unilist next-hop (e.g. an ae interface), mirrored packets may get dropped. [PR1134523](#)
- Kernel crash might be seen due to integer wrapping around in case of 64 bit architecture. [PR1134578](#)
- Insufficient time to allow an MPC5/MPC6 card to lock on the clocking source during FPC boot time might cause the Major Alarm raised due to "PLL Error." [PR1137577](#)
- MIC-3D-16CHE1-T1-CE only supports 4 queues by default due to the incorrect setting in code, this is a very minor change to make MIC-3D-16CHE1-T1-CE support 8 queues by default. [PR1138270](#)
- After removing a child link from AE bundle, the AE interface statistics in the SNMP MIB might show a spike. [PR1140533](#)
- When DHCP subscribers are brought up on the static interface IFL with interface-set, and this static interface IFL shares multiple DHCP stacks, it is possible that the interface-set does not get deleted when all DHCP subscriber are brought down on this static IFL. Unable to delete interface-set leads to commit denies on the dynamic profile involved. [PR1145450](#)
- Twice-NAT translation type does not work with the MS-MPC and MS-MIC service cards. The older MS-DPC cards does support this translation type. [PR1145690](#)
- With a 100G CFP2 MIC installed in a MPC6E FPC. If the FPC fails to initialize the MIC, it is very likely that the FPC will get into boot loop. [PR1148325](#)
- Subscriber traffic in an LNS coming from the core network is not switched properly when the incoming interface is an irb interface. [PR1148533](#)
- In EVPN environment, when CE MAC address alone gets changed for a MAC+IP entry, new MAC+IP entry is not getting reflected in EVPN database and the old entry still exists on PE router. [PR1149340](#)
- During deactivation of interfaces in a scaling setup the Packet Forwarding Engine may reboot or Packet Forwarding Engine may notice next-hop corruption. [PR1151844](#)
- From Junos OS release 14.2 with "exclude-hostname" configuration, hostname is not excluded from the messages before forwarding. This is a minor case, no other service impact. [PR1152254](#)
- Routers using inline layer 2 services may experience Packet Forwarding Engine wedge leading to fabric degradation and FPC restart. During issue state, the affected FPC will not be able to transmit and traffic will be fully blackholed. This problem is amplified by fragmented and out of order packets. This log entry may be seen during the error state: Host Loopback:HOST LOOPBACK WEDGE DETECTED IN PATH ID 0. [PR1153750](#)

- Copyright © 2018, Juniper Networks, Inc. 239

- In cases when the subscriber stacking is IPV6 over LNS, the IPV6 subscribers fails to come up with RPF check configured. DHC IPV6 subscriber over LNS comes up fine when RPF check configuration is disabled or removed. [PR1160370](#)
- Software OS thread on the line card is doing a busy loop by reading the clock directly from hardware. Sometimes it seems the thread is getting wrong values from HW register and waiting forever in the busy loop. After the busy loop crosses a certain time period, the line card crashes and reboots. This is a rare condition. [PR1160452](#)
- On MX Series routers with enhanced queuing DPCs, there is a memory leak whenever doing SNMP walk to any of COS related OID's or issue the command "show interfaces interface-set queue <interface set name>". [PR1160642](#)
- The Router Lifetime field is set to 0 in the first Routing Advertisement sent from LNS back to PPPoE subscriber. [PR1160821](#)
- The VCCPD_PROTOCOL_ADJDOWN system log message does not include a 'reason' string to explain why the virtual chassis adjacency was terminated. This information will now be present in the message. [PR1161089](#)
- When FPC goes to terminated state (FPC down, restarts) ACI interface-set does not get deleted. After FPC becomes online further subscriber bring up on this ACI interface-set fails. [PR1161810](#)
- Subscriber where TCP is attached to the underlying IFL will errantly end up in the control IFL queue. Workaround is to attach a TCP profile to each subscriber IFL. [PR1162108](#)
- Interfaces routing status message xxx.xxx.xxx.xxx <Up Broadcast> may be reported on an interface that is not associated with the config change, such as bridge-domain addition. It should be reported only if there is any change in the IFL parameters. This is an info(6) level message for debug purpose, so we can safely ignore the cosmetic problem. rpd[xxx]: %DAEMON-6: EVENT Flags ge-1/0/4.0 index 371 10.180.230.8/24 -> 10.180.230.255 <Up Broadcast> rpd[xxx]: %DAEMON-6: EVENT Flags irb.110 index 326 10.9.17.254/22 -> 10.9.17.255 <Up Broadcast> rpd[xxx]: %DAEMON-6: EVENT Flags irb.190 index 373 10.9.53.254/22 -> 10.9.53.255 <Up Broadcast> [PR1162699](#)
- MQCHIP reports continuous "FI Cell underflow at the state stage" message and continuous fabric drops on ADPC ICHIP Packet Forwarding Engines after ISSU on MX with ADPC. [PR1163776](#)
- The ability to configure a multicast group statically for a subscriber via a dynamic profile is not available in this release. Using the following statement, the subscriber can be enabled to receive multicast traffic for group 224.117.71.1 upon login: set dynamic-profiles <client profile> protocols igmp interface "\$junos-interface-name" static group 224.117.71.1 This support is not available and the subscriber needs to send a IGMP protocol JOIN message to receive multicast traffic. [PR1164323](#)
- On Junos OS Release 15.1 and later, on MS-MPC or MS-PIC, OSPF adjacency may fail to establish when there is no static route pointing to service PIC. [PR1164517](#)
- With IKEv1, MS-MPC packet drops on far-end after reboot of local MS-MPC. [PR1165787](#)
- When MS-MPC is used, if any bridging domain related configuration exists (e.g. "family bridge", "vlan-bridge", "family evpn", etc), in some cases, continuous MS-MPC crash hence traffic loss may occur. [PR1169508](#)

- PR1170019

- activated. [PR1170286](#)

- to forward about 62Gbps versus ~130Gbps and causing fabric queue drops. [PR1170805](#)

- automatically with fan at normal speed. [PR1174528](#)

- platforms.
- [PR1176575](#)

- Control Board.
- [PR1177571](#)

- MACSEC not working on layer 3 interface on MX104. [PR1177630](#)
- In a rare error scenario krt_q_entry of flow route was freed without dequeuing it from queue. This has been fixed via software change. [PR1178633](#)
- In MX Series running a Junos OS Subscriber Management Build, with more than 300+ firewall filters configured, it was found that an subscriber failed to login due to NACK received from system, stating the following error: BBE_DFW_DYN_PROF_ERR_STR session_id=1784: Can't find filter template named test300. BBE_DFW_DYN_PROF_ERR_CODE session_id=1784: Error code 13: Filter template not found. While the firewall filter named "test300" was certainly configured under the firewall filter configuration stanza; it found that the BBE daemon could hold a count of 256 filters only. Filters above this count were not getting indexed into the internal filter table and hence system could not find the filter. [PR1178671](#)
- In EVPN A/S mode, IFL mark down programming at the Packet Forwarding Engine on the BDF gets removed causing traffic loops. [PR1179026](#)
- [EVPN] Active-Active IP4 L3 session with CE over IRB Flaps. [PR1179105](#)
- When an MPC has training failure on all planes, then other MPCs in the system are getting affected. The root cause is that MQ MPC are not deleting the streams of the MPC which is causing the fabric wedge and effecting other MPCs. As a result FH is kicking in for other MPCs in the system. [PR1183230](#)
- When IPv4 firewall filter have 2625/32 destination in prefix-list, filter attached to subscriber interface is found broken. [PR1184543](#)
- Nexthop attribute in a framed route is not applicable anymore. Since subscriber IP address is used as the nexthop in all cases, there is no need to have an additional attribute for nexthop for framed routes. [PR1186046](#)

High Availability (HA) and Resiliency

- With NSR enabled on multiple Routing Engine system, when dynamic GRE tunnel is configured, performing Routing Engine switchover might cause rpd crash repeatedly on backup Routing Engine. [PR1130203](#)
- After graceful switchover is triggered in master VRRP router for the first time, the master state for all the VRRP instances are toggled to backup and comes back to master immediately. During this time all the traffic are dropped and comes back. [PR1142227](#)
- MXVC: ISSU failed after all FPC upgraded, TCP connection to kernel was dropped due to invalid IPC type 20. [PR1163807](#)

Infrastructure

- In scaling setup (in this case, there are 1000 VLANs, 1000 Bridge Domains, 120 IRB interfaces, 120 VRRP instances, BGP and IGP), if the routing protocols are deactivated and activated, there might be a chance that the pending route stats are not cleaned up, which will cause the stats infra to have stale pointers and lead to memory corruption in socket layers. The system might go to db prompt because of this. All the traffic going through the router will be dropped. [PR1146720](#)
- In Legacy, if the Routing Engine switchover on RPD crash configuration statement is enabled, the switchover is instigated in the Routing Engine kernel BEFORE the RPD core is created. In FreeBSD 10.x-based Junos OS, it is done AFTER. This creates an issue in scaled setups where the size of the RPD core, and therefore the time to create it, takes a lot longer. An FreeBSD 10.x-based Junos OS FreeBSD change has been made as part of this PR patch to adopt the same behavior as Legacy. [PR1164139](#)
- Support for fast Routing Engine mastership switchover when master Routing Engine crashes was broken in Occam releases up to and including 15.1R3. The patch for this PR addresses support for this feature in Occam in 15.1R4. [PR1167385](#)
- Issue is specific to Occam based images and is a very hard to induce. The issue occurs when virtual memory is low, and the FreeBSD OS tries to free-up memory by invoking the vm_lowmem event. In a multi-core environment, multiple kernel threads could encounter the low virtual memory condition at the same time, and so the vm_lowmem event could be invoked within the context of multiple kernel threads concurrently. Some of the protocols in the Junos networking stack register handlers against this event and two of these, clnp_drain() & tcp_drain(), were not SMP safe, which caused data corruption. clnp_drain() & tcp_drain() have now been made SMP safe; all other such handlers in the Junos networking stack were already SMP safe. [PR1182958](#)

Interfaces and Chassis

- Due to movement of SNMP stats model from synchronous requests to asynchronous requests in Junos OS Release 13.3R1, the IQ2/IQ2E PIC, which has limited memory and CPU power, can not handle scaling SNMP polling at high rate (e.g., a burst of 4800 SNMP requests). This issue comes with high rate SNMP stats polling for IQ2/IQ2E interfaces or Aggregated Ethernet (AE) interface with IQ2/IQ2E as member links. These memory failures can cause IQ2/IQ2E PIC reboot because keep alive messages will also not get memory. [PR1136702](#)
- When we polling SNMP MIBs for IPv6 traffic, for example, jnxIpv6IfInOctets, the logical interface (IFL) on IQ2 or IQ2E PIC may occasionally report double statistics. [PR1138493](#)
- %DAEMON-3-CHASSISD_I2C_WRITE_ERROR: i2cs_write_reg: write error for group 8 at address 0x49, offset 32 %DAEMON-3-CHASSISD_I2CS_READBACK_ERROR: Readback error from I2C slave for FPC 1 ([0x11, 0x42] -> 0x0) - The above errors represent transient communication issues between system components. - In certain cases, these can be service impacting. - Enhancements have been made for better handling of such error conditions. [PR1139920](#)

- On OAM maintenance domain intermediate Point (MIP), the connectivity fault management (CFM) will not be enabled on L2VPN interface if it is configured after L2VPN is up. [PR1145001](#)
- During a VRRP configuration change involving IP address change and/or VRRP configuration change while retaining same group ID, a race condition might occur causing vrrpd crash. [PR1145170](#)
- In affected releases, the following cosmetic alarms are seen after reseating the clocking cables: 2015-11-13 05:22:56 UTC Major CB 0 External-A LOS 2015-11-13 05:22:56 UTC Major CB 0 External-B LOS. [PR1152035](#)
- Remove MX Series from sending LCD halt message. [PR1153219](#)
- SONET interface on MIC-3D-IOC192-XFP does not count input error correctly. While hardware counts framing error, runts and giants but input error in 'show interface extensive' command reports runts and giant only. [PR1154268](#)
- When the master Routing Engine in the Virtual Chassis master router (VC-Mm) runs with high CPU (e.g. 99% CPU utilization), after a global/local switchover, the new master Routing Engine might relinquish its mastership during high CPU conditions. But the Virtual Chassis protocol role is not changed properly after the kernel relinquishes the mastership, causing dual master Routing Engines on this member router. [PR1156337](#)
- "monitor interface <if name>" will start ifmon process. In this time if telnet session to router is disconnected unconventionally, then ifmon process was not killed and it will take up 100% CPU utilization. The workaround is to terminate the stale ifmon process. [PR1162521](#)
- Internal timing for bringing FPCs online is extended for MX 2020/2010 systems to accommodate longer initialization times for fabric and FPCs. [PR1164147](#)
- CLI commit warning is replaced by syslog warning message when limited-ifl-scaling configuration statement is configured. Warning message text remains the same. [PR1165357](#)
- jpppd core at SessionDatabase::getAttribute() from Ppp::LinkInterfaceMsOper::getLowerInterfaceType() [PR1165543](#)
- If an interface configured with VRRP is removed from a routing-instance to global, or from global to a routing-instance, the Logical Interfaces of that interface will be deleted and recreated. In ideal case as the interface gets deleted VRRP should move to bringup state, when the interface is created again VRRP goes to previous state. After this VRRP should get VIP addition notification from kernel and update VRRP state and group id for VIP. However, in race conditions, VRRP might get VIP addition notification from kernel even before the interface creation event happens. If so VRRP will never be able to update proper VRRP state and group id. So the VIP will reply for the ARP with an incorrect MAC ending with "00" while the correct MAC should end with the groups id configured. [PR1169808](#)
- DCD core :/src/junos/sbin/dcd/infra/lag-link-dist/lag_link_dist_db.c:2147 [PR1175254](#)
- jpppd: RLIMIT_STACK & RLIMIT_SBSIZE messages are marked incorrectly at NOTICE level instead of at INFO level. [PR1178895](#)

- pppoe denies PADO for legitimate user PPPoE trace logs will report "Dropping PADI due to Duplicate Client" but there will be no subscriber logged in with that MAC address [PR1179931](#)
- Commit check may exit without providing correct error message and causing dcd exit. The only known scenario to trigger this issue is to configure a IPv6 host address with any other address on the same family. [PR1180426](#)

Layer 2 Features

- In BGP-based VPLS scenarios, changing the configuration of a VPLS mesh group might cause rpd core. FPC reboot might also be seen during the rpd core. [PR1123155](#)
- From Junos OS Release 13.2R1 and later, the rpd process might crash when adding/deleting Virtual private LAN service (VPLS) neighbors in a single commit. For example, a primary neighbor is changed to become the backup neighbor. [PR1151497](#)
- On GRES switch of mastership of Routing Engine via "request chassis routing-engine master switch", the dot1xd daemon will crash multiple times when 128K Logical Interfaces are configured in the MX960 chassis [PR1118475](#)
- On MX Series platforms, in DHCP subscriber management environment (the device is either used as local DHCP server or DHCP relay agent), if configuring the Aggregate Ethernet (AE) interface (e.g. change the "MTU" of AE) while there are subscribers on it, in race condition, the DHCP binding failure would occur on the AE. [PR1139394](#)
- In some cases where DHCP client devices are not fully protocol compliant they may become stuck trying to Renew an address lease indefinitely. These devices exposed a defect in the DHCP Relay behavior when acting as a proxy for the Server where a protocol NAK to restart the client was not properly created. As a result Address resources could be locked on the Relay preventing their use until the offending client device was restarted. [PR1153837](#)
- In Junos OS Release 15.1R3 with tomcat mode enabled, DHCP subscriber management with IRB interfaces is not reliable. It is possible that the DHCP bindings are unable to fully establish with IRB interfaces due to this reason. However, these bindings with same IRB interfaces should come up properly with tomcat disabled. [PR1155502](#)

MPLS

- In MPLS environment, the master Routing Engine might crash due to Mbuffer allocation failure and this crash will trigger an Routing Engine switchover, as a result Backup Routing Engine will become active. The issue is unreproducible, and trigger condition is not clear. [PR979448](#)
- During interoperation with CISCO device (e.g. CRS) belongs to different IGP area, if the P2MP LSP ping echo reply message from Cisco device is using interface address other than loopback/router-id as the source address, the reply message will be dropped on Junos OS device. With the fix, Junos OS device will accept the packets and print them as 'uncorrelated responses'. [PR1117166](#)
- Due to some data structure changes of ipc messages in 64-bit RPD, some of 32-bit applications (e.g. lsping, lspmon) would not work normally when RPD is running in

64-bit mode. Depends on Junos OS version, some of CLI commands might not work as expected. [PR1125266](#)

- While changing the label action for a static-label-switched-path from "stitch" to "pop", the routes added by stitch functionality is restored and there is no criteria for deleting the routes. Because of this, rpd crash might be seen. [PR1127348](#)
- MPLS TED might not select random links to calculate the ERO when OSPF is overloaded. Instead, only one or two interfaces will be used for all the configured LSPs originating from the router. [PR1147832](#)
- With RSVP refresh reduction feature enabled (using RSVP aggregate messages), when changing the configuration statement "no-load-balance-label-capability" to "load-balance-label-capability" on the egress router, the Entropy Label Capability (ELC) for the egress router would not be propagated towards the ingress. As a workaround, we can execute "clear rsvp session" on the ingress or wait until 3 refresh cycles (say 100s with default RSVP refresh config). [PR1150624](#)
- Static MPLS LSP using VT interface as a outgoing interface would not come up [PR1151737](#)
- LSPing returns 'routing instance does not exist' when used in vpls routing-instance under logical system. [PR1159588](#)
- If container LSP name and the suffix together are more than 60 characters in length, rpd process might crash during extensive split merge conditions. Its always advisable to keep them less than 60 characters. The member lsp name is coined in the following manner: <container name>-<suffix name>-<member count>- The LSP name can have upto 64 characters. So after putting together the container name, suffix, member-count (could go up to 2 digits), and the 2 hyphens, it should not exceed 64. So container-name and suffix together should not exceed 60 characters. A commit check will be added to throw warning if the name is more than supported character long. [PR1160093](#)
- When L2VPN composite next hop configuration statement is enabled along with L2VPN control-word, end-to-end communication fails. Because in this scenario, control-word is not inserted by the ingress PE, but other end expects the control-word. [PR1164584](#)
- Changing maximum-labels configuration under the hierarchy [edit interfaces interface-name unit logical-unit-number family mpls] might cause existing MPLS LSPs to become unusable. The root cause of this issue is that the family MPLS gets deleted and re-added. [PR1166470](#)
- In LDP-signaled VPLS environment, other vendor sends an Address Withdraw Message with FEC TLV but without MAC list TLV. The LDP expected that Address Withdraw Message with FEC TLV should always have MAC list TLV. As such, it rejected the message and close the LDP session. The following message can be seen when this issue occurs: A@lab> show log messages |match TLV RPD_LDP_SESSIONDOWN: LDP session xxx.xxx.xxx.xxx is down, reason: received bad TLV [PR1168849](#)
- In MVPN scenario, if active primary path goes down, then PLR(Point of Local Repair) needs to send Label Withdraw for old path and new Label Mapping for new path to the new upstream neighbor. In this case, LDP P2MP path may stay in "Inactive" state

for indefinite time if an LSR receives a Label Release, immediately followed by a Label Mapping for the same P2MP LSP from the downstream neighbor. [PR1170847](#)

- The rpd might crash upon receiving a TLE (Tag Label Element) delete notification arriving during a cleanup sequence. When adaptive teardown is configured and TLE delete notification comes during a cleanup sequence, this will trigger a recursive clean up and since the same cleanup routines are called and them being non-reentrant causes the code to assert. [PR1172567](#)
- When the egress LSR withdraws the label for its egress route, the rlfa nexthop for the ldp route for the egress remains in other routers running rlfs. A routing loop is formed when the rlfa nexthops for some of the router are pointing towards each other. Any traffic for the label route would loop until TTL expires. After the fix,rlfa nexthop with nexthop label alone will not be considered as valid lsp nexthop (primary nexthop). Ldp will send label withdraw for the label binding and delete the ldp route to avoid any potential routing loop. [PR1172581](#)

Network Management and Monitoring

- Eventd might run out of memory and crash because of excessive kernel logging. [PR1162722](#)

Platform and Infrastructure

- With "chassis maximum-ecmp 64" configured, when there is a route having 64 ECMP LSP next-hops and CoS-based forwarding (CBF) is enabled with 8 forwarding class ($64 \times 8 = 512$ next-hops), not all next-hops will be installed on Packet Forwarding Engine due to crossing the boundary in the kernel when number of ECMP next-hops is large than 309. [PR917732](#)
- When the 'enhanced-hash-key services-loadbalancing' feature is used by MX Series with MPCs/MICs based line cards, load balancing of flows across multiple service PICs via the source-address across does not work when internal BGP (IBGP) is used to steer traffic to the inside service-interface. For example the operator will see on the stateful firewall that the same source-address has flows across multiple service interfaces. [PR1034770](#)
- On MX Series-based platforms, when learning the MAC address from the pseudo-IFL (for example, label-switched interface), if the MAC address is aged out in the source FPC where the MAC got learned, due to the delay (around 2 to 3 milliseconds) of MAC address deleting message processed in the source FPC and the egress FPC (destination FPC of the traffic), the MAC address might be deleted first from the egress Packet Forwarding Engine but get added again during these 2-3 milliseconds time intervals (As there is continuous traffic coming on the egress FPC destined to this MAC, the MAC query is generated and sent to the Routing Engine and source FPC. Since the source FPC has not yet processed the MAC-deleted message, it sends the response, so stale MAC will get added on the egress Packet Forwarding Engine). In this situation, no L2 flooding would occur for the "unknown" unicast (since the MAC address is present on the egress Packet Forwarding Engine). [PR1081881](#)
- In certain cases, with some events such as disable/enable of links followed by Routing Engine rebooting or GRES enabled switch-over, below error message could be seen due to a software bug where it doesn't handle an internal flag properly. KERNEL/Packet

Forwarding Engine APP=NH OUT OF SYNC: error code 1 REASON: invalid NH add received for an already existing nh ERROR-SPECIFIC INFO: [PR1107170](#)

- Configuring one group with configuration of routing-instances and applying this group under routing-instances, then the rpd process will crash after executing "deactivating/activating routing-instances" commands. As a workaround, you can avoid using "apply-groups" under routing-instances hierarchy. [PR1109924](#)
- On MX Series with MPCs/MICs based linecard platform, if FPC offline is performed while FPC is in online progress (online process is at the stage of fabric links training), in very corner scenario, the Routing Engines state is stale and being sent to other existing FPCs, so the traffic forwarding might be affected. [PR1130440](#)
- Doing a file copy from a Routing-Engine running Junos OS image to a Routing-Engine running Junos OS with Upgraded FreeBSD image fails. [PR1132682](#)
- When there are additional messages related to FIPS generated during <commit configuration> rpc reply, the xml-tags closing tag <routing engine> may be missed in the reply. [PR1141911](#)
- FPC can crash and core due to a missing NULL check [PR1144381](#)
- During an ISSU upgrade in MXVC environment, linecards may crash causing service impact. When the linecards come up, there may be a nexthop programming issue as a secondary impact and some Logical Interfaces may not pass traffic. Affected linecards need to be rebooted to recover from this condition. [PR1152048](#)
- With Enhanced LAG mode enabled and sampling configured on AE interfaces, MS-DPC might drop all traffic as "regular discard". Disabling Enhanced LAG mode would avoid this issue. [PR1154394](#)
- On MX2000 Series platforms, when MPC goes down ungracefully, other MPCs in the chassis will experience "destination timeout". In this situation, auto fabric-healing will get triggered due to "destination timeout" condition, which may cause Fabric-Plane reset, even all other MPCs to be restarted in some cases. [PR1156069](#)
- cosd[20362]: cosd_config_database: Configuration database(/var/run/db/juniper-prop.data) does not exist. cosd[20460]: cosd_config_database: Configuration database(/var/run/db/juniper-prop.data) does not exist. The above log messages may be seen after after some commits. These messages do not pose an operational impact. [PR1158127](#)
- If one logging user is a remote TACACS/RADIUS user, this remote user will be mapped to a local user on device. For permissions authorization of flow-tap operations, when they are set on the local device without setting the permissions on the remote server, they cannot work correctly. The flow-tap operations are as follow: flow-tap -- Can view flow-tap configuration flow-tap-control -- Can modify flow-tap configuration flow-tap-operation -- Can tap flows [PR1159832](#)
- LU(or XL) and XM chip based linecard might go to wedge condition after receiving corrupted packets, and this might cause linecard rebooting. [PR1160079](#)
- NPC cored vpanic in
trinity_firewall_start_nh_get,trinity_firewall_add_and_check_internal,trini

ty_firewall_add_and_check. This line card core could potentially occur after an ISSU upgrade. [PR1160748](#)

- The following commit warning may be seen when using configure private and multi-line comments. This causes the commit to not complete. warning: outgoing comment does not match patch [PR1161566](#)
- Due to software bug on chassisd, backup CB temperature information is missing on cli command 'show chassis environment cb' if it's replaced once. [PR1163537](#)
- For MX Series Virtual Chassis with "default-address-selection" configured, when we have a discard route to a specific subnet (e.g. 10.0.0.0/8) with discard next-hop, and at the same time we have more specific routes through other interfaces (e.g. 10.1.1.1 through xe-0/0/0), if a UDP packet is being sent to 10.1.1.1 through xe-0/0/0 while interface xe-0/0/0 flaps or FPC reboots, it might cause kernel crash on both Master Routing Engine in the Virtual Chassis master router (VC-Mm) and Master Routing Engine in Virtual Chassis backup router (VC-Bm). As a workaround, we can disable "default-address-selection" configuration. [PR1163706](#)
- Below log can be seen on MX2020 after One FPC was pulled out and committing the configuration related interface. CHASSISD_UNSUPPORTED_FPC: FPC with I2C ID of 0x0 is not supported [PR1164512](#)
- A sonet interface configured as unnumbered BFD session fails to come up. [PR1165720](#)
- Modifying the configuration of a hierarchical policer when in use by more than 4000 subscribers on an FPC can cause the FPC to core and restart. [PR1166123](#)
- There are three issues related to DDOS reported in the PR 1168425. 1) Some policers are configurable, but do not react when disabling them (tunnel-ka aggregate, re-services-v6 capti.v6, syslog aggregate) With the fix all the configurable DDOS protocol parameter changes will get reflected correctly in Packet Forwarding Engine. 2) Some policers for non-unclassified traffic are non-configurable (control aggregate, mcast-snoop mld, ipsec aggregate, uncls resolve-v4, uncls resolve-v6, uncls filter-v4, uncls filter-v6, tunnel-ka aggregate). These policers are internally deprecated or renamed and not shown on CLI anymore. So any configuration will not come to the Packet Forwarding Engine sides. 3) Some policers are for unclassified traffic are non-zero (mlp unclass, services unclass, radius unclass, ip-frag unclass, gre unclass, re-services unclass, re-services-v6 unclass) We do not have a convention of setting unclassified to 0. Consider this as FAD. [PR1168425](#)
- In Junos OS Release 15.1, a customized password prompt that can be sent by a TACACS+ server is not displayed to the user upon login. A usual password prompt "Password: " is displayed instead. The issue is seen when the following conditions are met: 1. Junos OS Release 15.1 without the fix for this PR is used. 2. TACACS+ is used for the user authentication 3. When user logs in, TACACS+ server sends a customized password prompt for this user. For example, this can cause an issue when S/KEY-based one-time password (OTP) authentication is configured for a particular user on the TACACS+ server because the user might be unable to calculate the one-time password as they would not see the key sequence number and the seed provided by the authentication server. [PR1168634](#)

- Because the sequence number in RPM ICMP-PING probes is introduced as 32-bit variable instead of 16-bit, if it increases and reaches the max value 65535, it does not rollover, which might cause all RPM ICMP-PING probes to fail and not succeed any more. [PR1168874](#)
- In affected release, if user runs the Packet Forwarding Engine debug command like "show sample-rr eg-table ipv4 entry ifl-index 1224 gateway 113.197.15.66" will cause the MPC crash. [PR1169370](#)
- Long container elements can have keys which could be very big in size. If the key is more than 256, max key length in Patricia tree, mustd is coring, which leads router into amnesiac mode and any login is denied. [PR1169516](#)
- Layer 2 protocols might flap when router was flooded with low priority traffic reaching towards FPC CPU/Routing Engine CPU when DDoS protection is disabled. [PR1172409](#)
- On MPC5E/6E/7E/8E/9E/NG linecards, firewall filter of family inet/inet6/vpls configured with non-contiguous prefixes for address matching might fail and cause traffic drop. Using only contiguous prefixes can avoid this issue. [PR1172725](#)
- On all Junos OS platforms, when using RADIUS server, after RADIUS request is successfully sent by Junos device, if the network goes down suddenly, then response sent by the RADIUS server is not received within timeout period. In this scenario, the RADIUS request will be sent again with invalid socket descriptor, which will lead to auditd (provides an intermediary for sending audit records to RADIUS and/or TACACS+ servers) crash. [PR1173018](#)
- "show arp" command can't get complete results and reports "error: could not find interface entry for given index". [PR1174150](#)
- On MX2010/2020, MPC/SFB cards do not boot up if single phase AC PSMs are turned ON sequentially with interval even though the PSMs have sufficient remaining power. [PR1176533](#)
- A flow is determined by doing hashing on the packet header. Usually 5-tuple (src/dest IP addresses, IP protocol number, src/dest ports) are used for hashing because a flow is defined by 5-tuple. This is all fine for TCP/UDP packets. But layer-3 packets generated by JDSU tester only have layer-3 header and don't have layer-4 header. JDSU tester uses the same location as layer-4 header as packets' sequence number. So MX Series with MPCs/MICs card treats sequence number of JDSU tester packets as layer-4 header of a packet, hence, Junos thinks every packet is a single flow and order of different flows are not guaranteed. [PR1177418](#)
- JFLOW: when IPv6 route point to AE bundle, jflow record shows Outgoing Interface as child interface and not actual AE interface. [PR1177790](#)

Routing Policy and Firewall Filters

- interface-routes rib-group import-policy is not in effect to filter prefixes correctly. All direct prefixes could be installed into the secondary route table. [PR1171451](#)

Routing Protocols

- When configuring router in RR mode (cluster-id or option B MP-eBGP peering), the advertise-external feature will not be applicable in local VRFs due to a different route selection/advertisement process (main bgp.l3vpn.0 vs VRF.inet.0). [PR1023693](#)
- BFD session configured with authentication of algorithm keyed-sha1 and keyed-md5 might be flapping occasionally due to FPC internal clock skew. [PR1113744](#)
- During many types of configuration changes, especially including import policy, BGP has the need to re-evaluate the routes it has learned from peers impacted by the configuration change. This re-evaluation involves re-running import policy to see if there is any changes to the learned routes after applying the new policy. This work is done in the background as part of an "Import Evaluation" job. When BGP is reconfigured a second time, and the "Import Evaluation job" has not completed, it is necessary to re-run the job from the beginning if there's another change to policy or something with similar impact. This state is noted as "Import Evaluation Pending". However, in this case, there was a bug that caused BGP to always enter the pending state upon reconfiguration, regardless of whether relevant changes were made to import or other similarly impactful configuration. The result is that once it is necessary to start re-evaluation of the routes for a peer, even trivial configuration changes that happen too quickly will cause the "Import Evaluation job" to need to run again as a result of the "Pending" flag being set. To avoid the issue, please ensuring that "ImportEval" is not present in a BGP peer's Flags output from the CLI (show bgp neighbor) prior to doing even trivial commits. [PR1120190](#)
- In multicast environment, when the RP is FHR (first hop router) and it has MSDP peers, when the rpf interface on RP changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)
- In a situation which BGP is being used in combination with interface's rfp-check; deleted routes may see delay in propagation of BGP withdrawn messages. [PR1135223](#)
- When interface IP MTU is less than 1464 bytes and the number of LSA headers in an OSPF DbD packet is big enough for it to exceed the MTU (i.e. OSPF database contains enough LSAs), unexpected fragmentation of OSPF DbD packets may occur due to incorrect calculation of maximum allowed payload size. [PR1148526](#)
- In BGP scenario with large scale routing-instances and BGP peers configured, due to a software defect (a long thread issue), BGP slow convergence might be seen. For example, BGP might go down 8-9 seconds after BFD brings down the EBGP session. The rpd slip usually does not hurt anything functionally, but if the slip gets big enough, it could eventually cause tasks to not be done in time. For example, BGP keepalives with lower than 90 seconds hold-time might be impacted. There is no known workaround for this issue, but configuring configuration statement "protocol bgp precision-timers" can take care of the weak spot like sending BGP Keepalives. [PR1157655](#)

- Starting from Junos OS Release 15.1R1 to Junos OS Release 15.1R3, and Junos OS Release 15.1F2 to Junos OS Release 15.1F4, Junos OS devices may not be able to establish BGP sessions with legacy router that does not support BGP optional parameters. The reason is that capability of supporting BGP open message fallback to no optional parameter is removed in these releases, which causes "OPEN Message Error (2)" during session setup. [PR1163245](#)
- In BGP scenario with independent domain enabled in a VRF, when configuring a BGP session in a VRF routing instance with a wrong local-as number, some routes might be declared as hidden because of AS path loop. If later configuring the correct AS number as local-as and committing the configuration, those routes might still remain in hidden state. The hidden routes can be released after performing commands "commit full" or "clear bgp table <ANY_VRF>.net.0". [PR1165301](#)
- In L3VPN scenario, feature multipath is configured under [set protocols bgp group] with L3VPN chained CNH under routing-options, the feature multipath does not work for L3VPN routes. [PR1169289](#)
- When clearing IS-IS database, process rpd might crash due to a rare memory de-allocation failure that a task pointer is attempted to be freed twice. In the fix of this issue, the order of referencing the task pointer is being revised to avoid the occurrence of rpd crash. [PR1169903](#)
- PIM bootstrap export policy is not working as expected when there are no pim neighbors up on the router [PR1173607](#)

Services Applications

- When making a configuration change to a EXP type rewrite-rule applied to a SONET interface in an MX FPC Type 2 or MX FPC Type 3, if MS-DPC is also installed on the device, a MS-PIC core dump may be generated. [PR1137941](#)
- In a rare situation in a SIP conversation we might end up in a situation where we have a child conversation whose entry is still present in the parent conversation while the child flow is already deleted. While trying to delete this child flow from the parent conversation validate if the flow is valid and go ahead with deleting the child flow. [PR1140496](#)
- When deleting NAT flow under a race condition the Service PIC can core [PR1159028](#)
- These log messages no longer appear in syslog if log level is set to warning / error or higher. If the log level is set to notice or lower (info / debug) then these log messages are shown in syslog file. [PR1162116](#)
- In Layer 2 Tunneling Protocol (L2TP) subscriber management environment, the jl2tpd process (L2TP daemon) might crash during clean-up of L2TP tunnel or session after it failed to establish. [PR1162445](#)
- When traffic is flowing through MS-DPC card Service PIC and there is an active port block and some ports are assigned from that active port block, if changing the max-blocks-per-address setting to a lower value (lower than the current value), the service line card may crash. [PR1169314](#)

- MS-PIC core-dump when MPLS or IPV6 routing updates are received. This is a race condition rarely seen while IPV6 or MPLS routes are deleted or added in the MS-PIC. [PR1170869](#)
- Attempting to ping a subscriber address from the L2TP LNS CLI will fail. [PR1187449](#)

Subscriber Access Management

- The range for the request-rate statement at the [edit access radius-options] hierarchy level has been extended to 100 through 4000 requests per second. In earlier releases, the range is 500 through 4000 requests per second. The default value is unchanged at 500 requests per second. [PR1033668](#)
- If a DHCP local pool is exhausted, the newly dialed in subscriber B might get the IP address of newly logged out subscriber A, in a very rare condition, if the acc-stop message for A is sent to Radius server after acct-start for B, and if the Radius server identify the subscribers only by IP address but not by session, the subscriber B might get terminated. [PR1079674](#)
- In DHCP relay scenario, DHCP relay binding might get stuck in "RELEASE(RELAY_STATE_WAIT_AUTH_REQ_RELEASE" state due to the LOGOUT Request is not processed correctly by authentication manager process (authd) if there were multiple attempts to activate Lawful Intercept (LI) for this DHCP subscriber using RADIUS change of authorization (CoA) packets in quick succession. [PR1179199](#)

User Interface and Configuration

- From Junos OS Release 13.2R1 and later, the commitd process might crash while committing large configurations in a single commit, for example, committing 250k lines of config on top of existing config. This issue is due to a lack of storage space for current and running configurations. [PR1159462](#)

VPNs

- Upon clearing p2mp lsp in dual-home topology, system is adding the same outgoing interface to the (S,G)OIL multiple times and thus duplicate/multiply the amount outgoing traffic. [PR1147947](#)

Resolved Issues: 15.1R3

- [Class of Service \(CoS\) on page 254](#)
- [Forwarding and Sampling on page 254](#)
- [General Routing on page 255](#)
- [High Availability \(HA\) and Resiliency on page 268](#)
- [Infrastructure on page 268](#)
- [Interfaces and Chassis on page 269](#)
- [Layer 2 Features on page 273](#)
- [MPLS on page 275](#)
- [Network Management and Monitoring on page 276](#)

- [Platform and Infrastructure on page 277](#)
- [Routing Protocols on page 282](#)
- [Routing Policy and Firewall Filters on page 284](#)
- [Services Applications on page 284](#)
- [Software Installation and Upgrade on page 285](#)
- [Subscriber Management and Services on page 285](#)
- [User Interface and Configuration on page 288](#)
- [VPNs on page 288](#)

Class of Service (CoS)

- The chassis-scheduler-map is not applied to interface if FPC restart, Routing Engine switchover, or reboot. Only after deactivation/activation of the affected interface does the CoS get applied again. [PR1132983](#)
- When the system has "system services subscriber-management enable" set (means the subscribers are VBF flow based), the ICMP MTU exceed notification may not be sent to subscribers, which will cause the subscriber Path MTU Discovery to fail. [PR1138131](#)
- This PR does optimization in AE SNMP handling. If all the links in an AE bundle go down, then any CoS SNMP query for this AE IFD/IFL will return cached values. [PR1140440](#)
- On the MX104 platform, when applying the "rate-limit" and the "buffer-size" on the logical tunnel (lt-) interface on the missing MIC (not inserted on MPC), commit failure with error message would occur. As a workaround, this issue could be avoided by applying the "rate-limit and "buffer-size" on inserted MIC, then commit. [PR1142182](#)

Forwarding and Sampling

- The command "clear firewall all" will now clear the policer stats displayed by "show policer __auto_policer_template_1__", ... "show policer __auto_policer_template_8__". [PR1072305](#)
- This issue is seen in Junos OS Release 14.2 and later releases. When Routing Engine based sampling is enabled and BGP session is using 4 byte AS, improper AS number can be found in sampling information. [router1]-----[DUT]-----[router2] AS 1,000 A AS 10,0000 | sampling 1.1.1.1 ----->2.2.2.2 traffic --- traceoptions log --- Aug 10 12:21:21 v5 flow entry Aug 10 12:21:21 Src addr: 1.1.1.1 Aug 10 12:21:21 Dst addr: 2.2.2.2 Aug 10 12:21:21 Nhop addr: 20.20.20.1 Aug 10 12:21:21 Input interface: 747 Aug 10 12:21:21 Output interface: 749 Aug 10 12:21:21 Pkts in flow: 594 Aug 10 12:21:21 Bytes in flow: 49896 Aug 10 12:21:21 Start time of flow: 4648545 Aug 10 12:21:21 End time of flow: 4707547 Aug 10 12:21:21 Src port: 0 Aug 10 12:21:21 Dst port: 2048 Aug 10 12:21:21 TCP flags: 0x0 Aug 10 12:21:21 IP proto num: 1 Aug 10 12:21:21 TOS: 0x0 Aug 10 12:21:21 Src AS: 1000 Aug 10 12:21:21 Dst AS: 34464 <<<<< Aug 10 12:21:21 Src netmask len: 32 Aug 10 12:21:21 Dst netmask len: 32. [PR1111731](#)
- On the MX Series platform with MX-FPC/DPC, M7/10i with Enhance-FEB, M120, M320 with E3-FPC, when there are large sized IPv6 firewall filters(for example, use prefix lists with 64k prefixes each) enabled, commit/commit check would fail and the dfwd process would crash after configuration commit/commit check. There is no operational impact. [PR1120633](#)

- On all Junos OS platforms, when both the filter and the policer are configured for an interface, in rare cases, the policer template may not be received by the Packet Forwarding Engine (from the Routing Engine) when it is referenced by the filter term (normally the policer template gets received before the filter term referencing it which is ensured by mechanism in the Routing Engine kernel). In this situation, the FPC would crash due to this rare timing issue. This issue might be avoided by the recommended steps below: 1. Deactivate the physical interface (IFD) and commit 2. Enable any filter and policer that attached to the interface (e.g. IFL) and commit 3. Activate interface back. [PR1128518](#)
- On MX80 and MX104 platform, applying firewall filter with MX Series specific match condition will raise the following warning message. Filter <filter_name> is MX Series specific; will not get installed on DPCs for interface <interface_name>. This warning message is needed for the other modular type MX Series platforms since it can have DPC and MPC mixed. But the message is not needed for MX80 and MX104 platform since they only have the MX Series based Packet Forwarding Engine. Although the warning message tells that the relevant firewall filter is not installed, the firewall filter is correctly installed into Packet Forwarding Engine. Thus, user can ignore the message in case the warning message is logged on MX80 and MX104 platform. [PR1138220](#)
- For Junos OS release 14.1R1 and later, when a broadcast packet is sent in a scenario of Integrated routing and bridging (IRB) over Virtual Tunnel End Point (VTEP) over IRB, the packet is getting dropped in kernel as it was looping due to a software issue. The error log message "if_pfe_vtep_ttp_output: if_pfe_ttp_output failed with error 50" is observed when issue occurs. [PR1145358](#)
- On MX Series-based platforms, in race condition, when using the policer which has configuration statement "bandwidth-percent" configured (e.g., set firewall policer XXX if-exceeding bandwidth-percent 80), if the logical interface (IFL) bandwidth change and the filter bind message arrive at the Packet Forwarding Engine out of order (e.g., when changing the bandwidth of the IFL or rebooting the FPC), the "bandwidth-percent" policer may end up using physical interface (IFD) bandwidth for "bandwidth-percent" computation. [PR1154034](#)

General Routing

- On an MX Series Virtual Chassis platform, when we restart one or both of the standby Routing Engines, the log message "ksyncd_select_control_plane_proto: rhost_sysctlbyname_get: No such file or directory" might be observed as the ksyncd daemon attempts to select a communication protocol (UDP/TCP). After several tries, it will fall back to TCP and proceed as normal. [PR945925](#)
- In an MX Series Virtual Chassis (MX-VC) environment, the private local nexthops and routes pointing to private local next hops are sent to the Packet Forwarding Engine from the master Routing Engine and not sent to the slave Routing Engine, then a Routing Engine switchover happens. Now as the new master Routing Engine does not know about such next hops and routes, they are not cleaned up. When a next hop with same index is added on the new master Routing Engine and sent to the Packet Forwarding Engine, the Packet Forwarding Engine might crash due to a stale next hop exist. [PR951420](#)

- In a Layer 3 wholesale configuration, DHCPv6 advertise messages might be sent out with source MAC all zeroes if the subscriber is terminated on the demux interface in a non-default routing instance. For subscribers on default instance there is no such issue observed. [PR972603](#)
- On MX Series routers with MPC3E, MPC4E, MPC5E, and MPC6E, Junos OS does not support short(sub-second) interface hold-time down configuration. So, a hidden configuration statement is introduced to ignore DFE tuning state during hold-down timer period. This configuration statement allows sub-second hold-down timer on MPC3E,MPC4E,MPC5E,MPC6E. set interfaces <intf name> hold-time up <U ms> down <D ms> alternative The configuration statement does not work/support 'MPC5E 3D Q 2CGE+4XGE' and 'MIC6 2X100GE CFP2 OTN', and we recommend configuring hold-time down to be more than 3 seconds for these two cards. [PR1012365](#)
- On MX240/480/960/2010/2020 platform with Junos OS release 15.1R1 and later, the process health monitor process (pmond) is not available on the Routing Engine. The msppmond process on MS-MIC/MS-MPC tries to connect pmond process on Routing Engine continuously but fails. It will result in additional traffic between the MS-MIC/MS-MPC and Routing Engine, causing high CPU utilization. [PR1014584](#)
- No performance or functional impact. Can be safely ignored. "Ignore the PTP message (2) as this MPC doesn't support EEC" should be moved from notice to debug level. [PR1020161](#)
- MIC-3D-8OC3-2OC12-ATM Revision 22 or later is supported only by the following Junos OS releases: Junos OS Release 12.3 — 12.3R9 and later, Junos OS Release 13.3 — 13.3R6 and later, Junos OS Release 14.1 — 14.1R4 and later, Junos OS Release 14.2 — 14.2R3 and later, Junos OS Release 15.1 and later. [PR1036071](#)
- There is a remote loop back feature in 802.3ah standard, where one end can put remote end into remote-loopback mode by sending enable loopback control lfm PDU. In remote loopback, all incoming packets (except lfm packets) are sent back on wire as it is. Transmit or receive of lfm packets should not be affected when an interface is in remote loopback mode. On the VMX platform when we configure the lfm remote-loopback we run into problem state, In problem state we will see that LFM packets sent from node which is in loopback state is not reaching the peer end hence we will not see the remote entity information for the "run show oam ethernet link-fault-management" command on peer router. [PR1046423](#)
- On all routing platforms M Series, MX Series, T Series with BGP configured to carry flow-specification route, in case of deleting a filter term and policer, then add the same term and policer back (it usually happens in race condition when adding/deleting/adding the flow routes), since confirmation from dfwd for the deleting policer might not be received before attempting to add the same policer, the rpd would skip sending an add operation for it to dfwd. As a result, when the filter term is sent to dfwd and tell it to attach to the policer, dfwd had already deleted the policer, and since rpd skipped re-adding it, dfwd will reject the attach filter with policer not found error and rpd will crash correspondingly. [PR1052887](#)
- When a labeled BGP route resolves over a route with MPLS label (e.g. LDP/RSVP routes), after clearing the LDP/RSVP routes, in the short window before the LDP/RSVP

routes restore, if the BGP routes resolves over a direct route (e.g. a one-hop LSP), the rpd process might crash. [PR1063796](#)

- [illegible]

```
Temperature Intake 36 degrees C / 96 degrees F FPC 9 status: State Online
Temperature Intake 37 degrees C / 98 degrees F FPC 10 status: State Online
<<<<<<<<<<<<<<<<<<< Correct info Temperature Intake 42 degrees C / 107 degrees
F {master} > show chassis fpc detail Slot 0 information: State Online Temperature 36
<<<<<<<<<<<<<<<<<<< Correct info Total CPU DRAM 3584 MB Total XR2 518 MB Total
DDR DRAM 49920 MB Start time: 2015-05-12 12:36:14 AST Uptime: 9 days, 1 hour, 31
minutes, 38 seconds Max Power Consumption 1088 Watts Slot 9 information: State
Online Temperature 37 Total CPU DRAM 2048 MB Total RLDRAM 331 MB Total DDR
DRAM 1280 MB Start time: 2015-05-12 12:38:00 AST Uptime: 9 days, 1 hour, 29 minutes,
52 seconds Max Power Consumption 239 Watts Slot 10 information: State Online
Temperature 42 <<<<<<<<<<<<<<<<<<< Correct info Total CPU DRAM 3584 MB Total
XR2 518 MB Total DDR DRAM 49920 MB Start time: 2015-05-12 12:36:18 AST Uptime:
9 days, 1 hour, 31 minutes, 34 seconds Max Power Consumption 1088 Watts PR1090671
```

- When wrong diagnostic optics info might be seen for GE-LX10 SFP and SFP+ for SumitomoElectric. The issue occurs for a specific SFP type - "Xcvr vendor part number : SCP6F44-J3-ANE", it can be seen with "show chassis pic fpc-slot X pic-slot Y".

```
user@device> show chassis pic fpc-slot 0 pic-slot 0 .. PIC port information: Fiber Xcvr
vendor Wave- Xcvr Port Cable type type Xcvr vendor part number length Firmware 0
GIGE 1000LX10 SM OPNEXT INC TRF5736AALB227 1310 nm 0.0 1 GIGE 1000LX10 SM
FINISAR CORP. FTLF1318P2BTL-J1 1310 nm 0.0 2 GIGE 1000LX10 SM SumitomoElectric
SCP6F44-J3-ANE 1310 nm 0.0 <<<<Error SFP>>>>PR1091063
```
- Occasionally, AFEB PCI reads from Cortona MIC with ATM OAM traffic might return garbage values even though the actual content in the MIC has the correct value, this corrupted values would lead to AFEB crash, and also PCI error logs such as :

```
afeb0 PCI ERROR: 0:0:0:0 Timestamp 91614 msec. afeb0 PCI ERROR: 0:0:0:0 (0x0006)
Status : 0x00004010 afeb0 PCI ERROR: 0:0:0:0 (0x001e) Secondary bus status :
0x00004000 afeb0 PCI ERROR: 0:0:0:0 (0x005e) Link status : 0x00000011 afeb0
PCI ERROR: 0:0:0:0 (0x0130) Root error status : 0x00000054 afeb0 PCI ERROR:
0:0:0:0 (0x0134) Error source ID: 0x02580258 afeb0 PCI ERROR: 0:2:11:0 Timestamp
91614 msec. afeb0 PCI ERROR: 0:2:11:0 (0x0006) Status : 0x00004010 afeb0 PCI
ERROR: 0:2:11:0 (0x004a) Device status : 0x00000004 afeb0 PCI ERROR: 0:2:11:0
(0x0052) Link status: 0x00004001 afeb0 PCI ERROR: 0:2:11:0 (0x0104) Uncorrectable
error status: 0x00000020 afeb0 PCI ERROR: 0:2:11:0 (0x0118) Advanced error cap
& ctl : 0x0000001e5 afeb0 PCI ERROR: 0:2:11:0 (0x011c) Header log 0 : 0x00000000
afeb0 PCI ERROR: 0:2:11:0 (0x0120) Header log 1 : 0x00000000 afeb0 PCI ERROR:
0:2:11:0 (0x0124) Header log 2 : 0x00000000 afeb0 PCI ERROR: 0:2:11:0 (0x0128)
Header log 3 : 0x00000000 PR1097424
```
- When the clock sync process (clksyncd) is stopped and resumed during link flaps, the clksyncd process might get into an inconsistent state with various symptoms, the clock source might be ineligible due to "Interface unit missing" or "Unsupported interface" with no Ethernet Synchronization Message Channel (ESMC) transmit interfaces.
[PR1098902](#)
- In abnormal session close scenario like by pulling-out running ms-mpc or in scaled flow environments, some garbage object can remain due to a bug on internal flow state machine then would trigger mspmand coredump. The fix of this PR clears such a problematic status objects. [PR1100363](#)

- After Junos OS Release 13.3R1, IPCMON infra is added to debug IPCs between PFEMAN and the Routing Engine. When convergence occurs, string processing of IPCMOM will take added time. Then the slow convergence will be seen. It is a performance issue, it is visible in scaled scenario (for example, more than 100K routes). As a workaround, please execute command "set pfe ipclog filter clear" to disable IPC logging on all FPCs. [PR1100851](#)
- Fragmenting a special host outbound IP packet with invalid IP header length (IP header length is greater than actual memory buffer packet header length), can trigger NULL mbuf accessing and dereferencing, which may lead to a kernel panic. [PR1102044](#)
- On MX Series platforms, in subscriber management environment, when carrying scaling subscribers, as the Packet Forwarding Engine process (pfed) memory usage will grow along with the number of subscribers, the pfed memory usage limit may get reached (that is, 512M) because of the subscriber scale and number of service attached to the subscribers (for example, when carrying more than 140k single stack PPPoE subscribers per chassis, 4 services per subscriber), in this situation, the pfed crash may occur due to memory exhaustion. [PR1102522](#)
- On MX Series platform, in subscriber management environment, if the subscriber's underlying logical interface (IFL) is static (for example, ge-x/y/z.0 or aex.0 rather than ge-x/y/z.32767 or aex.32767) with family inet configured, when all the subscribers are logged out, the ARP on the underlying IFL may stop resolving the next-hop path due to the incorrect deletion of ARP family of the underlying IFL when removing subscribers. [PR1102681](#)
- With Nonstop active routing (NSR) enabled, deleting routing-instance/logical system configuration might cause a soft assert of rpd. If NSR is not enabled, after deleting routing-instance/logical system configuration, executing "restart routing" might trigger this issue too. The core files could be seen by executing CLI command "show system core-dumps". This timing issue has no function impact. [PR1102767](#)
- cpcdd core observed in scaled scenario. [PR1103675](#)
- On MX Series platform, when using DHCPv6 Prefix Delegation (DHCPv6-PD) and DHCPv6 subscriber sessions are layered over an underlying PPPoE subscriber session, due to the fact that the value of the UDP checksum in Echo reply message might get incorrectly set to all zero (i.e. "0x0000"), a small number (for example, on a 1 to 5 subscribers out of 10000 subscribers basis) of subscribers might fail to renew the IPv6 addresses in each lease time circle. [PR1103349](#)
- When using "write coredump" to invoke a live coredump on an FPC in T Series, the contents of R/SR ASIC memory (Jtree SRAM) will get dumped. In the situation that there is a parity error present in the SRAM, then the coredump will abort and the FPC will crash. As a workaround, configuring "set chassis pfe-debug flag disable-asic-sram-dump" before "write coredump" will help to avoid the issue. [PR1105721](#)
- When mspmand (which manages the Multiservice PIC) core dump (when the mspmand crash, it will dump a core file for analysis) is in progress in MS-MPC/MS-MIC and a GRES command is issued at the same time, it is seen that the MS PIC gets stuck and has to be recovered by offlining/onlining the PIC. [PR1105773](#)

- Dynamic vlan ifl is not removed with 'remove when-no-subscriber' configuration. [PR1106776](#)
- When Bridge domain in PBB-EVPN Routing instance is modified to add/remove ISIDs BD can get stuck in destroyed state. This happens when ISIDs in the Bridge domain are changed from 1 to many or many to 1. This is only noticed during configuration changes or initial deployment. [PR1107625](#)
- Under IPv6 VRRP scenario, when a host sends router solicitation messages to VRRP virtual IPv6 address, the VRRP master replies router advertisement messages with physical MAC address instead of virtual MAC, the VRRP slave replies router advertisement messages with physical MAC address as well. As a result, the host has two default gateways installed and the host will send traffic directly to two devices but not to the VRRP virtual IP. This issue affects VRRP function and traffic. [PR1108366](#)
- On MX Series platform with "subscriber-management" enabled, while high scaled subscribers (for example, 126K dual-stack DHCP v4/v6 subscribers over VLAN demux) login/logout at high rate, MX Series-based line cards which hold subscribers might crash after the bbe-smgd process restart. [PR1109280](#)
- On MX240/480/960 Series router with MS-DPC, customer running BGP over IPsec. This BGP session has a BFD session tied to it. The BGP session is up but the BFD session remains in INIT state. The issue might be seen with any service configured with multihop BFD enabled. Traffic forwarding will not be affected. [PR1109660](#)
- In subscriber management environment and the accessing interface is an AE interface, after AE interface flap or FPC reboot, the subscriber traffic accounting might not be reported on demux interface but on the underlying AE interface. [PR1110493](#)
- In rare condition, after Routing Engine switchover, the MPC PIC might offline, and some error messages might be seen. [PR1110590](#)
- This issue is a regression defect introduced in Junos OS Release 11.4R11, 12.1R10, 12.2R8, 12.3R6, 13.2R4, 13.3R2, 14.1R1. After upgrading to those releases containing the original fix, when there is no export policy configured for forwarding table to select a specific LSP, whenever routes are resolved over RSVP (for example, due to aggressive auto-bandwidth), resolver will spend considerable amount of time on resolver tree, which contributes to base line increase in rpd/Routing Engine CPU. [PR1110854](#)
- Resolved problem with Syslog messages generated like "krt_decode_resolve for 239.255.255.250, 101.11.67.33: no logical interface for index 1073741825" when Multicast packets are received on Subscriber interfaces. [PR1110967](#)
- On MX Series platform, when using FTP Application-level gateway (ALG), if the FTP (including both active mode and passive mode) server requests client to use different IP address for control session and data session (i.e. after the control session is established, the destination IP address of FTP server is changed on which client should transfer the data), although the control session could be built, the data session could not be established due to wrong pinhole creation. The issue would not occur in the scenario that the port is changed while the destination IP address is the same. [PR1111542](#)
- CLI core dump is due to repeated mismatched XML open/close directives in the "show pppoe lockout" output. This issue is most likely to occur when there is a ratio of 8 PPPoE clients in lockout per VLAN. [PR1112326](#)

- Right now this fix is available from Junos OS Release 14.2R6 and later. On Junos OS Release 14.2R5 or older images MSRPC gates once opened would never get deleted. From Junos OS Release 14.2R6 and later, MSRPC gates are opened for 60 mins no matter whether expected packet hits gate or not. After 60 minutes gates are deleted by timer. [PR1112520](#)
- In the scenario that the power get removed from the MS-MPC, but Routing Engine is still online (for example, on MX960 platform with high capacity power supplies which split into two separate power zones, when the power zone for the MS-MPC line card loses power by switch off the PEM that supports the MS-MPC situated slot), if the power goes back (for example, switch on the PEM), the MS-MPC might be seen as "Unresponsive" (checked via CLI command "show chassis fpc") and not coming up back online due to failure of reading memory. [PR1112716](#)
- Under certain conditions, when the Junos OS Routing Engine tries to send an IP packet over a IPIP tunnel, the lookup might end up in an infinite loop between two IPIP tunnels. This is caused by a routing loop causing the tunnel destination for Tunnel#A to be learned through Tunnel#B and the other way round. [PR1112724](#)
- On all Junos OS platform, when the Junos Routing Engine tries to send an IP traffic over a GRE tunnel, the route lookup might end up in an infinite loop between two GRE tunnels (the infinite loop is caused by a routing loop causing the tunnel destination for Tunnel A to be learned through Tunnel B and the other way round), the kernel would crash as a result. As a workaround, the issue could be avoided by preventing the tunnel destination of a tunnel to be learned through a second tunnel (and the other way round). [PR1113754](#)
- On MX Series Virtual Chassis with "subscriber-management" enabled, the BBE subscriber management daemon (bbe-smgd) might crash on the Virtual Chassis backup router (VC-Bm) during subscribers concurrent login/logout. The bbe-smgd process restart results in a minimal impact of system and subscribers. All connected subscribers remain connected. [PR1113792](#)
- On MX Series routers with Junos OS release 12.3X54-D20 or 12.3X54-D25, Inverse multiplexing for ATM (IMA) interfaces on MIC-3D-4COC3-1COC12-CE may not come up due to "Insufficient Links FE" alarm. This is due to data corruption on the physical layer. [PR1114095](#)
- On MX-VC with heartbeat connection, if it is in a scaled subscribers environment, when power down both VCM Routing Engines, there might be a delay (minutes) for backup chassis to be master and during which time, traffic blackhole might be seen. [PR1115026](#)
- On MX Series Virtual Chassis with "subscriber-management" enabled, the BBE subscriber management daemon (bbe-smgd) might crash on the Virtual Chassis backup router (VC-Bm) during subscribers concurrent subscribers churn. The bbe-smgd process restart results in a minimal impact of system and subscribers. All connected subscribers remain connected. [PR1115187](#)
- After VC Protocol Master Switch, new VCMm could allocate STP index of 1 (which is global discarding state) to new Physical Interfaces resulting in STP status incorrectly marked to discarding on the FPCs of the current VCBm. Please note for the fix to be effective, it is required that MXVC setup is rebooted once after upgrade of all the Routing

Engines of the MXVC chassis with new fixed image following normal upgrade procedure and hence ISSU based upgrades are not supported. [PR1115677](#)

- On a busy MX Series Virtual Chassis platform, for example, with 100k subscribers and 16k subscribers concurrent login/logout, the ksyncd process might crash on Virtual Chassis backup Routing Engines after a local or global graceful Routing Engine switchover (GRES). This issue has no service impact. [PR1115922](#)
- For MPC6E with CFP2, there was a race condition between the Interrupt service routine and the periodic, as a result interface up/down will not happen for laser off/on. [PR1115989](#)
- On MX240/MX480/MX960 platform with MS-DPC card, in some race conditions, after deactivating member interface of the aggregated multiservices (AMS) interface, the service PIC daemon (spd) might crash due to memory corruption. As a workaround, we should offline the member PICs before changing the AMS configuration and then online the PICs. [PR1117218](#)
- On M Series /MX Series platform, the 10G Tunable SFP/SFP+ can not be tuned in Junos OS Release 15.1R2. [PR1117242](#)
- In broadband edge (BBE) environments with graceful Routing Engine switchover (GRES) enabled, the BBE subscriber management daemon (bbe-smgd) might crash on the master Routing Engine after Routing Engine switchover. [PR1117414](#)
- On MX Series routers containing multiple Packet Forwarding Engines such as MX240/MX480/MX960/MX2010/MX2020, with either MPC3E/MPC4E/MPC5E/MPC6E cards, if the routers have GRE decap, then certain packet sizes coming via these aforementioned line cards, at very high rate can cause these line cards to exhibit a lockup, and one or more of their Packet Forwarding Engines corrupt traffic towards the router fabric. [PR1117665](#)
- During the LSP switch-over, the hiwatermark may get set to unexpectedly high value. The issue happens due to incorrect reference point taken while calculating the Max avg BW in the last interval and this results in incorrect Highest Watermark BW in the autobadnwidth stats. [PR1118573](#)
- alg-logs and pcp-logs are not supported under [edit edit services service-set <ss name> syslog host local class] on ms interface as of now. Added warning message for the same during configuration commit. [PR1118900](#)
- On MX Series platform, in rare condition, if removing or deactivating "member-interfaces" configured for an aggregated Multiservices (AMS) bundle (only officially supported on MS-MPC/MS-MIC), for example, using CLI command "deactivate interfaces ams0 load-balancing-options member-interface mams-7/1/0", all the MX Series-based FPCs and the MS-MPC/MS-MIC may crash. As a workaround, to avoid the issue, below is the recommended procedures to change AMS bundle size, 1. Offline member PICs 2. Change AMS configuration 3. Online member PICs [PR1119092](#)
- The rpd process might crash when executing CLI command "show evpn database" with the combination of "vlan-id" and "mac-address". [PR1119301](#)
- In the multicast environment with pd interface (interface on the rendezvous point (RP) that de-encapsulates packets), if execute GRES multiple times, and the GRES interval

is less than 30 minutes, the routes on master Kernel are added and deleted for a short while. In rare condition, backup Kernel will not be able to see them. So after Routing Engine switchover, the new master Kernel will delete next-hop ID for such routes, but Packet Forwarding Engines will not see this deleted message. As a result, the Kernel/Packet Forwarding Engine are out of sync for such particular next-hop ID, it might trigger a reset of all the Packet Forwarding Engines. As a workaround, please do the Routing Engine switchover more than 30-minute intervals. [PR1119836](#)

- On MS-MPC equipped MX Series platform, during the "three-way handshake" process, when receiving ACKs (e.g. after sending SYN and receiving SYN/ACK) with window size 0 (as reported, it is set to 0 by TCP client when using some proprietary protocol), the ACKs would be incorrectly dropped by the line card due to failure in TCP check. This issue could be avoided by preventing software from dropping packets that fail in the check, for example, by CLI command below, re# set interfaces ms-3/0/0 services-options ignore-errors tcp. [PR1120079](#)
- The commands "show igmp interface <interface name>" and "show mld interface <interface name>" may sometimes result in memory corruption and cause a core dump of smg-service daemon. [PR1120484](#)
- The commit latency will increase along with the increasing lines under [edit system services static-subscribers group <group name> interface]. Use ranges to create static demux interfaces is a recommended option. e.g.: [edit system services static-subscribers group PROFILE-STATIC_INTERFACE] + interface demux0.10001001 upto demux0.10003000; [PR1121876](#)
- On MX240/MX480/MX960/MX2010/MX2020 products with MPC2E-3D-NG/MPC2E-3D-NG-Q/MPC3E-3D-NG/MPC3E-3D-NG-Q with MIC-3D-4XGE-XFP, Physical Interface flap detection is much slower. It might lead to high FRR time, some traffic might be lost. [PR1122589](#)
- ovs-vxlan -- irb mac address is missing in ovs database. [PR1122826](#)
- For scaled configuration, it may take too much time for commit, and session gets hung because there is an unnecessary check to see if family Ethernet-switching co-exists with family bridge for all interfaces having bridge configuration. [PR1122863](#)
- MX Series router acting as L2TP access concentrator (LAC) may not recognize the MLPPP protocol field (0x003d) in the inbound PPP packet from customer premise equipment (CPE) and could disconnect the session not respecting idle-timeout. The traffic forwarding might be affected. [PR1123233](#)
- When MX-VC is under a high latency transport condition (usually happens in DDoS attack), the performance might reduce and the backup Routing Engine's unnecessary and harmful resync operations could ultimately consume the entire available /mfs buffer space, which finally resulting in traffic loss. [PR1123842](#)
- On MX Series platform, the MS-MPC crash may occur. The exact trigger of the issue is unknown, normally, this issue may happen over long hours (e.g. within a week) of traffic run (e.g. running HTTP/HTTPS/DNS/RTSP/TFP/FTP traffic profile). [PR1124466](#)
- With BGP configured on CE-faced interfaces (in VRFs), doing 'show route' frequently may cause rpd to slowly leak memory. The leak rate will be one memory block of the size necessary to hold the instance name of the routing instance for a BGP neighbor.

If the rpd process memory gets exhausted, the rpd process might crash, and the routing protocols are impacted and traffic disruption will be seen due to loss of routing information. You can check rpd memory usage with "show task memory brief" command. [PR1124923](#)

- Right now this fix is available from Junos OS Release 14.2R6 and later. On Junos OS Release 14.2R5 or older images SUN RPC gates once opened would never get deleted. From Junos OS Release 14.2R6 and later, SUN RPC gates are opened for 60 minutes no matter whether expected packet hits gate or not. After 60 minutes gates are deleted by timer. [PR1125690](#)
- In multihoming EVPN scenario and the customer facing interface is an AE interface, after moving an interface from the EVPN instance into a VPLS instance, traffic loss might be seen on CE facing FPC. [PR1126155](#)
- In EVPN scenario, the EVPN route table between the master Routing Engine and backup Routing Engine would be different (unused garbage routes will appear) once Routing Engine switchover (e.g., by rebooting the "old" master Routing Engine or performing graceful routing engines switchover) is performed, which may cause kernel crash on the new master Routing Engine in some cases. [PR1126195](#)
- When Junos OS devices use Link Layer Discovery (LLDP) Protocol, the command 'show lldp neighbors' displays the contents of PortID Type, Length, and Value (TLV) received from the peer in the field 'Port Info', and it could be the neighbor's port identifier or port description. Junos OS CLI configuration statement can select which 'interface-name' or 'SNMP ifIndex' to generate for the PortID TLV, so we do not have any problem as long as two Junos OS devices are connected for LLDP, but we might have an interoperability issue if other vendor device which can map the configured 'port description' in the PortID TLV is used. In such case, Junos OS displays the neighbor's PortDescription TLV in the 'Port info' field, and if the peer sets 'port description' whose TLV length is longer than 33 byte(included), Junos is not able to accept the LLDP packets then discards packets as errors. The PortID TLV is given as : "the port id tlv length = port description field length + port id subtype(1B)". [PR1126680](#)
- EVPN route attributes like the label and Ethernet segment identifier (ESI) may be missing from EVPN family routes installed by BGP. [PR1126770](#)
- In multi-homing Ethernet VPN (EVPN), if there are two loopback addresses and the router-id and the primary loopback addresses are different on the designated forwarder (DF) PE, when the link between CE and DF PE down, the Type 4 route of old DF are not deleted properly from the backup PE and causing the new DF election failure. The traffic forwarding will be affected. As a workaround, we should configure single primary loopback address and remove "router-id" configuration statement on both multi-homing PEs. [PR1126875](#)
- On M320/T320/T640 with FPC 1/2/3 and their enhanced version (-E2/-E), in multicast scenario and AE interface is within multicast NH (such as, AE interface is the downstream interface for a multicast flow), egress multicast statistics displays incorrectly after flapping of AE member links. [PR1126956](#)
- An incorrect destination MAC address is applied to the packet when a DHCPv6 Offer/Advertise packet is sent back to the subscriber from a non-default routing instance across a pseudowire. [PR1127364](#)

- On MX Series platform with "subscriber-management" enabled, when a dynamic DHCPv4 subscriber is stacked over a static VLAN and the "route-suppression access-internal" configuration statement is enabled, before the subscriber is established, it is possible for ARP process to first add a resolved route matching the subscriber's IP address. Then when the subscriber is established, the subscriber management process will change this route, but the change is not handled properly in the Packet Forwarding Engine. Due to this timing issue, the broadband network gateway (BNG) fails to forward transit packets to this subscriber. For example, the external DNS server's response packets might not be delivered to the voice subscriber interface resulting in voice service outage. As a workaround, we can disable "route-suppression". [PR1128375](#)
- On MX Series platform, when offlining the line card (possibly, with any of the line cards listed below), "Major alarm" might be seen due to HSL (link between line card and Packet Forwarding Engine) faults. This fault is non-fatal and would not cause service impact. The line cards that may hit the issue could be seen as below, MS-MPC/MS-MIC MIC-3D-8DS3-E3 MIC-3D-8CHDS3-E3-B MIC-3D-4OC3OC12-1OC48 MIC-3D-8OC3OC12-4OC48 MIC-3D-4CHOC3-2CHOC12 MIC-3D-8CHOC3-4CHOC12 MIC-3D-1OC192-XFP MIC-3D-1CHOC48. [PR1128592](#)
- In current Juniper implementation, the IPv6 multicast Router Advertisement timer is not uniformly distributed value between MinRtrAdvInterval and MaxRtrAdvInterval as described in RFC 4861. [PR1130329](#)
- When software encounters an error configuring the optics type into the VSC8248 PHY retimer component of an MX MIC/PIC (typically done on SFP+ module plugin), this could lead to 100% FPC CPU utilization indefinitely. MPCs and MICs that are potentially affected are: MPC3 + 10x10GE SFPP MIC MPC4 32XGE MPC4 2CGE+8XGE (10G interfaces only) MPC6 + 24x10GE (non-OTN) SFPP MIC. [PR1130659](#)
- On MX with MS-MIC (or possibly, MS-MPC is affected as well), changing configuration of sampling input parameters, such as "rate" under forwarding-options is not reflected without restarting the line card. [PR1131227](#)
- On MX Series based line cards, multiple modifications of firewall filter might cause lookup chip error and traffic blackhole, following jnh_free error messages could help to identify this issue: messages: fpc1 jnh_free(10212): ERROR [FW/3]:1 Paddr 0x006566a9, addr 0x2566a9, part_type 0 call_stack 0x40497574 0x418ffa84 0x41900028 0x418ecf94 0x41861690. [PR1131828](#)
- CLI output of "clear services sessions" gives an impression to the user that session is marked for deletion in case of delayed delete but the XML output "clear services sessions|display xml" of the above command says "session removed". Ideally both should convey the same message to the user. The changes have been made to make sure CLI and XML information given to the user in sync. [PR1132006](#)
- Packet logs were not available in previous releases. Now in X55-D35 onwards and in mainline from (exact 14.2, 15.1 releases numbers to be determined), these logs will be available.. [PR1132162](#)
- When customers do changes under "protocol router-advertisement interface X" (such as changing timers etc), they expect that commit would trigger a new router-advertisement being sent out to notify hosts about configuration changes.

However it does not seem to be a case unfortunately. It makes the router information to expire on hosts and causes obvious loss of connectivity for the hosts. [PR1132345](#)

- In subscriber management environment with autosense VLAN, if IP demux interface is not configured, the IGMP/MLD join message from client might be dropped due to "Bad Receive If". [PR1132929](#)
- The subscribers login rate could be degraded when IGMP/MLD is enabled on the dynamic demux interface. [PR1134558](#)
- On MX Series platforms with non-Q MPC (for example, MPC2-3D) or Q-MPC with enhanced-queueing off, when traffic has to egress on any one of the dynamic PPPoE (pp0), IP-DEMUX (demux0) and VLAN-DEMUX (demux0) Logical Interfaces, the queue mapping might get wrong. The traffic forwarding might be affected. [PR1135862](#)
- While bringing down subscribers, the system generates [Deinstantiate Service Failed permanently, daemon: cosd] error message. [PR1136083](#)
- MXVC-Same subnet VC-heartbeat polling failed to recover. [PR1136119](#)
- On MX Series platforms with MIC3-3D-1X100GE-CFP, after In-Service Software Upgrade (ISSU), the Junos upgrade is successful, but the 100GE port will be down, and the traffic forwarding will be affected. [PR1136269](#)
- In IGMP over subscriber environment with configuration statement "remove-when-no-subscribers" configured, after performing graceful Routing Engine switchover, subscribers with multicast joins cannot re-login when subscriber logout before it sends IGMP leave in new master. [PR1136646](#)
- On MS-MIC, TCP session Up/Down causes JSERVICES_NAT_* and JSERVICES_SESSION_* messages though severity level "none" is configured for services. [PR1137596](#)
- JNH periodically attempts to recover memory no longer in use. Recently when Firewall address space was expanded to 16M, a side effect was triggered -- memory recovery was extended to 16M as well. On the Hercules line card, Firewall does not use a small block of IDMEM, causing JNH to attempt the return of the unused memory. There is no mechanism for recovery of IDMEM, therefore, this message is displayed. Excepting the syslog impact, there is no further effect on the line card. [PR1140021](#)
- From Junos OS Release 14.1R4, 14.2R3, 15.1 and later, when firewall filter is applied to NG-MPC, after system reboot, Routing Engine might go into amnesiac mode. [PR1141101](#)
- In subscriber management environment, on MX Series platform, after login/logout static subscribers (e.g. by setting/deleting the interface), some of the static subscribers may get stuck in "Terminated" state. [PR1143205](#)
- When multicast-only fast reroute (MoFRR) is enabled in PIM or multipoint LDP domain, memory leak will be observed on generation of the multicast FRR next-hops. The leak rate is 8-byte for IPv4 and 12-byte for IPv6 addresses, per FRR next-hop created. Eventually, the rpd process will run out of memory and crash when it cannot honor some request for a memory allocation. [PR1144385](#)

- When ARP is trying to receive a nexthop message whose size (for example 73900 bytes) is bigger than its entire socket receive buffer (65536 bytes), the kernel might crash, and the traffic forwarding might be affected. [PR1145920](#)
- On MX Series routers with "subscriber-management" enabled, the BBE subscriber management daemon (bbe-smgd) might crash on the backup Routing Engine when performing graceful Routing Engine switchover (GRES) during subscribers concurrent login/logout. [PR1147498](#)
- On MX Series platform, in multicast subscriber management environment (e.g. IGMP is configured for subscribers in dynamic profile), when nonstop active routing (NSR) is enabled, if the routing protocol process (rpd) is busy or there are hundreds of multicast groups are active (e.g., 250), missing multicast entries issue might be seen after performing Routing Engine switchover twice or more (i.e., first Routing Engine switchover works fine, and the issue may occur from the second switchover and onward). As a workaround, this issue could be avoided by issuing CLI command "restart smg-service" on backup Routing Engine after every switchover. [PR1149065](#)
- When a routing instance is configured with "routing-instances <instance name> routing-options localized-fib" then VPN localization may fail, causing all routes for the affected routing instance to be installed on all Packet Forwarding Engines. [PR1149840](#)
- Commit error after attempting to delete all guaranteed rates on all traffic-control-profiles associated with demux0 [edit] lab@mx480-J12_09# commit re0: [edit class-of-service interfaces] 'demux0' IFL excess rate not allowed on interface (demux0), please specify guaranteed rate on at least one IFL error: configuration check-out failed. [PR1150156](#)
- When using type 5 FPC on T4000 platform, traffic go out of the interface where "source-class-usage output" is configured will be dropped if the Source class usage (SCU) or Destination Class Usage (DCU) policy configuration is missing. This issue is caused by incomplete configuration so, to avoid the issue, please make the configuration complete (e.g. with "source-class-usage output" and SCU policy). [PR1151503](#)
- In the TXP environment, the Line-Card Chassis (LCC) Switch Interface Board (SIB) status is not right when execute command "user@router> show chassis environment", their status are Absent, but no alarms. This is a minor issue, it does not affect business. [PR1156841](#)
- A previous enhancement to strengthen the VC-Heartbeat message exchange resulted rejecting messages at the crucial time of determining the health of the other VC member when all adjacency links fail. Validation of messages has been adjusted to remain strong when the VC is connected, but relaxed during the split conditions to prevent rejecting valid messages. [PR1157383](#)
- In PPPoEv6 scenario, the unsolicited Router Advertisement will be sent out before get IPCPv6 ack. This behavior will impact PPPoEv6 connection rate. We can use "no-unsolicited-ra" configuration statement to suppress this message as a workaround. But in this case, this configuration statement does not work. The unsolicited Router Advertisement will still be sent out. [PR1158476](#)

High Availability (HA) and Resiliency

- On MX Series platforms with Junos OS Release 15.1R1 and later, while a core dump is in progress, if we try to access the dump directory, due to the deadlock defect, the system might hang and crash. As a workaround, we should not access the "/var/crash" directory till the core dump is complete. [PR1087082](#)
- On MX Series Virtual Chassis (MX-VC) with scaled configuration, for example, 110000 DHCP and 11600 PPP subscribers, the unified in-service software upgrade (ISSU) might fail due to the management daemon (MGD) timer expiring before Field-replaceable units (FRUs) update finish. [PR1121826](#)
- On MX240/480/960/2010/2020 platform with Junos OS Release 15.1R1 and later, in high scale scenario (e.g., there are 4 million routes or more), the connection between Routing Engine and the FPC(s) may flap after performing graceful Routing Engine switchover (GRES). The other symptoms are intermittent packet drops between the Routing Engine and FPC during regular operation without performing GRES and scaled scenario. [PR1146548](#)

Infrastructure

- Only the following directories and files are preserved when upgrading from build prior to 15.1 to 15.1 (FreeBSD 10) . config/ /etc/localtime /var/db/ /var/etc/master.passwd /var/etc/inetd.conf /var/etc/pam.conf /var/etc/resolv.conf /var/etc/syslog.conf /var/etc/localtime /var/etc/exports /var/etc/extensions.allow /var/preserve/ /var/tmp/baseline-config.conf /var/tmp/preinstall_boot_loader.conf Anything else not listed above is deleted/formatted during upgrading to freebsd10 version of Junos OS. [PR959012](#)
- When "show version detail" CLI command has been executed, it will call a separate gstatd process with parameter "-vvX". Because the gstatd could not recognize these parameters, it will run once without any parameter then exit. In result of "show version detail", following information could be seen: user@hostA> show version detail
Hostname: hostA Model: mx960 Junos: 13.3R6-S3 JUNOS Base OS boot [13.3R6-S3]
JUNOS Base OS Software Suite [13.3R6-S3] .. <snipped> file: illegal option -- v usage:
gstatd [-N] gstatd: illegal option -- v usage: gstatd [-N] <snipped> At the same time, log lines like following might be recorded in syslog: file: gstatd is starting. file: re-initializing gstatd mgd[14304]: UI_CHILD_START: Starting child '/usr/sbin/gstatd'
gstatd: gstatd is starting. gstatd: re-initializing gstatd gstatd: Monitoring ad2 gstatd: switchover enabled gstatd: read threshold = 1000.00 gstatd: write threshold = 1000.00
gstatd: sampling interval = 1 gstatd: averaged over = 30 mx960 mgd[14304]:
UI_CHILD_STATUS: Cleanup child '/usr/sbin/gstatd', PID 14363, status 0x4000
mgd[14304]: UI_CHILD_EXITED: Child exited: PID 14363, status 64, command '/usr/sbin/gstatd' [PR1078702](#)
- On dual Routing Engine platforms, if GRES is configured (triggered by "on-disk-failure"), when a disk I/O failure occurs on the master Routing Engine due to hardware issue (for example, SSD failure), the graceful Routing Engine switchover might not be triggered immediately after initial IO failure has been detected. As a result, the Routing Engine might enter a state in which it responds to local pings and interfaces remain up, but no other processes are responding. [PR1102978](#)

- With scaled configuration or there are memory leaks, if the virtual memory is running very low, the kernel might crash and the device will go in db prompt continuously due to a recursion issue. [PR1117548](#)
- The "show route vpn-localization" command does not have any output, but if xml format requested then xml output of the same command works. [PR1125280](#)
- In scaling setup (in this case, there are 1000 VLANs, 1000 Bridge Domains, 120 IRB interfaces, 120 VRRP instances, BGP and IGP), if the routing protocols are deactivated and activated, there might be a chance that the pending route stats are not cleaned up, which will cause the stats infra to have stale pointers and lead to memory corruption in socket layers. The system might go to db prompt because of this. All the traffic goes through the router will be dropped. [PR1146720](#)

Interfaces and Chassis

- Reconfiguring lt- interface causes dcd memory leak. [PR 879949](#)
- On MX Series routers, the physical or logical interfaces (ifd/iftl) might be created and marked UP before a resetting FPCs' fabric planes are brought up and ready to forward traffic, as a result, traffic might be black-holed during the time window. This window of traffic black-hole is particular long if the chassis is heavily populated with line-cards, for example, the router has large scale of configuration (routes or subscribers), and coupled with a lot of FPC reset, such as upon a node power up/reset. [PR918324](#)
- jnxBoxDescr is reworded for MXVC to replace the platform type with a more general representation that replaces the specific member platform type with "Virtual Chassis". Old virtual chassis text example: jnxBoxDescr.0 = member0 Juniper MX240 Internet Backbone Router New virtual chassis text example: jnxBoxDescr.0 = member0 Juniper MX Virtual Chassis Internet Backbone Router NOTE: The MIB design for jnxBoxAnatomy "top-level" chassis information works properly for a standalone chassis, but doesn't fully represent virtual chassis multi-member configurations because it is capable of providing information for only one physical chassis. (The remainder of the jnxBoxAnatomy MIB "containers" properly support the inventory of a multi-member configuration.) MX virtual chassis provides another MIB, jnxVirtualChassisMemberTable, to supply the equivalent "top-level" information. [PR1024660](#)
- When issuing a CFM LTR from CE, link state reply, recieved from MX Series, acting as MHF doesn't contain Reply Egress TLV if ingress and igress IFL are located on the same Physical Interface [PR1044589](#)
- MS-DPC might crash when allocating chain-composite nexthop in enhanced LAG scenario. [PR1058699](#)
- During subscriber login/logout the below error log might occur on the device configured with GRES/NSR. /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV (pp0.1073751222). [PR1058958](#)
- Currently the redundant logical tunnel (rlt) interface only supports limited vlan range (0..1023), it should support the extended vlan range (0..4094) as the logical tunnel does. [PR1085565](#)
- Trap messages does not logged on logical interface (iftl) after deleting "no-traps" configuration statement, in spite of setting explicit "traps". [PR1087913](#)

- The Enhanced LAG feature is enable in network-service enhanced-ip mode, but it is not supported in enhanced-ethernet mode. [PR1087982](#)
- During scaling login/logout different types of subscribers (e.g. 17K) on LAC router, there might be some L2TP LAC subscribers stuck in terminating state and never get cleared, blocking new sessions from establishing on the same interface. [PR1094470](#)
- When DHCP subscribers are terminated at specific routing-instances and the interface stack is IP demux over vlan-subinterface over AE interface, there might be a memory leak in kernel AE iffamily when subscribers login/logout. [PR1097824](#)
- The adaptive load balancing counters are always zero for aggregated Ethernet (AE) bundles on MICs or MPCs of MX Series routers. [PR1101257](#)
- VRRP inet6 group interface does not send Router Advertisement (RA) when the interface address and virtual address are same. run show ipv6 router-advertisement interface ge-0/2/0.430 Interface: ge-0/2/0.430 Advertisements sent: 0 Solicits received: 0 Advertisements received: 0 [PR1101685](#)
- With "enhanced-ip" mode and AE interface configured, if SCU/DCU accounting is enabled, the MS-DPC might drop all traffic as regular discard. [PR1103669](#)
- The 'optics' option will now display data for VCP ports: show interfaces diagnostics optics vcp-0/0/0 [PR1106105](#)
- On MX240 or MX480 platform with at least two DC modules (PN: 740-027736) equipped, when shutting down one of the PEMs and then turn it on again, even the PEM is functioning, the "PEM Fan Fail" alarm might be observed on the device due to software logic bug. There is no way to clear the ALARM_REASON_PS_FAN_FAIL for I2C_ID_ENH_CALYPSO_DC_PEM once it has been raised. [PR1106998](#)
- On MPC-3D-16XGE-SFPP line card, when an optics (for example, 10G-LR-SFP) is disabled and then enabled administratively, if the SFP is not temperature tolerant (non-NEBS compliant), the TX laser may not be turned on due to the fact that the chassis process (chassisd) may keep sending the "disable-non-nebs-optics" command to the optics if the current temperature of FPC reaches the threshold temperature. [PR1107242](#)
- On MX Series platforms, continuous error messages might be seen on the MICs (for 10G/40G/100G MICs) from MIC3 onwards (listed as below) when physical interface (IFD) settings are pushed (e.g. booting the MPC). Based on the current observation, the issue may not have any operational impact and the MICs that may encounter this issue are listed as below, - 10G MICs: MIC3-3D-10XGE-SFPP, MIC6-10G, MIC6-10G-OTN, - 40G MICs: MIC3-3D-2X40GE-QSFPP, - 100G MICs: MIC3-3D-1X100GE-CFP, MIC3-3D-1X100GE-CXP, MIC6-100G-CXP, MIC6-100G-CFP2 [PR1108769](#)
- Junos OS now checks ifl information under the ae interface and prints only if it is part of it [PR1114110](#)
- The jpppd process (which is used to authenticate subscribers) might crash after restarting MPC in live network, and then some subscribers might be found stuck in INIT state. [PR1114851](#)
- In PPPoE subscriber management environment, when dynamic VLAN subscriber interfaces is created based on Agent Circuit Identifier (ACI) Information, the subscribers

might unable to login after reboot FPC with syslog "Dropping PADI due to no ACI IFLSET". [PR1117070](#)

- When an M120/M320/MX Series router acts as the Broadband network gateway (BNG) and provide the PPPoE subscriber management service, after Routing Engine switchover, it might wrongly send out IPCP Term-Req message. It will cause PPPoE subscribers login failure. [PR1117213](#)
- When using Ethernet OAM Connectivity Fault Management (CFM), the CFM process (CFMD) may crash in either of the following scenarios, - Scenario 1 When CFMD is restarted or GRES. There is no specific defined configuration which could cause this crash, but normally this would be seen with VPLS or Bridge domain with multiple Mesh-groups. The crash happens rarely in this scenario. - Scenario 2 When configuring 2 interfaces in the same bridge-domain (BD) or routing-instance, and both interfaces have maintenance association end point (MEP) configuration along with action-profile enabled. Also there is no maintenance association intermediate point (MIP) configuration on that BD or routing-instance. The crash might be seen with the above configurations and when one of the interfaces is flapped or deleted and then re-created. In addition, in this scenario, this issue may not happen always as this depends on the ordering of kernel event. [PR1120387](#)
- The jpppd process might crash and restart due to a stale memory reference. The jpppd process restart results in a minimal impact of system and subscribers. All connected subscribers remain connected and only subscribers are attempting to connect at time of process restart would need to retry. [PR1121326](#)
- On Junos OS platforms, an aggregate-ethernet bundle having more-than one member link can show incorrect speed which would not match to the total aggregate bandwidth of all member links. The issue would be seen when LFM is enabled on the aggregate-ethernet bundle. The issue would be triggered when one of the member link flaps. Although after the flap, the current master Routing Engine would show correct aggregate speed, the backup Routing Engine would report incorrect value. In this state, when Routing Engine mastership is switched, the new master Routing Engine (which was backup) will show incorrect value. One of the side-effect of this issue is that RSVP also reflects incorrect bandwidth availability for the affected aggregate-ethernet bundle, thus can cause under-utilization of the link with LSP having bandwidth constraints. [PR1121631](#)
- Since a bug which was introduced in Junos OS Release 15.1R1, loopback sub-interfaces always have a Flag down in the output of CLI command "show interfaces". [PR1123618](#)
- The connectivity fault management (CFM) log message "Adjacency up" should only be logged when the router first detects remote MEP or the peer interface goes down and up causing adjacency failure for this remote MEP. But now it is wrongly logged when any peer set/clear the Remote defect indication (RDI) bit in continuity check messages (CCMs). [PR1125164](#)
- If two redundant logical tunnels (rlt) sub-interfaces are configured in a same subnet and in a same routing-instance, a sub-interface will be down (this is expected), but if the sub-interface is removed from the routing-instance later, after disable and enable the rlt interface, a sub-interface might remain in down state unless removing configuration of rlt interface and then rollback. [PR1127200](#)

- With incomplete cfmd configuration, for example, only MD (maintenance-domain) configured and no MA (maintenance-association) configured, or MD and MA configured but no MEP configured, SNMP walk in CFM MD table results in infinite loop and process cfmd is spinning at around 90% CPU. [PR1129652](#)
- In Dynamic PPPoE subscriber management scenario, when the system is overloaded with requests coming, the subscribers might fail to login in a race condition. [PR1130546](#)
- The jpppd process might crash and restart due to a buffer overwrite. The jpppd process restart results in a minimal impact of system and subscribers. All connected subscribers remain connected and only subscribers are attempting to connect at time of process restart would need to retry. [PR1132373](#)
- MX-VC specific behavior for SNMP walk of jnxOperating* containers was divergent from physical MX. Returned to vergence. [PR1136414](#)
- On MX Series platforms, the "Max Power Consumption" of MPC Type 1 3D (model number: MX-MPC1-3D) would exceed the default value due to software issue. For example, the value might be shown as 368 Watts instead of 239 Watts when "max ambient temperature" is 55 degrees Celsius. [PR1137925](#)
- When Micro Bidirectional Forwarding Detection (BFD) sessions are configured for link aggregation group (LAG), the device control process (DCD) acts as the client to the micro BFD session. In order to monitor the connection between client (DCD) and server(BFD), client needs to exchange keep alive hello packets with the server. To send hello packets, DCD needs to move out of IDLE phase to CONFIG_BFD phase which is the reason for below log messages: dcd.c:585 dcd_new_phase_if_idle() INFO : Current phase is IDLE, going to phase CONFIG_BFD usage.c:75 dcd_trace_times() INFO : Phase Usage for IDLE : user 0.001 s, sys 0.000 s, wall 60.019 s dcd.c:717 dcd_new_phase() INFO : New phase is CONFIG_BFD usage.c:75 dcd_trace_times() INFO : Phase Usage for CONFIG_BFD : user 0.000 s, sys 0.000 s, wall 0.000 s dcd.c:717 dcd_new_phase() INFO : New phase is IDLE There is no functionality impact, however these messages may flood the logs. As a workaround, we can filter out these messages from being written to the log file according to this [KB article](#). [PR1144093](#)
- In MX-VC or VRR platforms running releases of 15.1 built before about February 2016, the following cosmetic warning message will be print upon commit: [edit] 'chassis' warning: WARNING: MPC reboot or chassis reboot is required to use MIC aware dynamic power management feature on already plugged-in MPCs. [PR1144295](#)
- The alarm "CB 0 ESW Packet Forwarding Engine Some Ports Failed " was triggered by the difference "rcb_handle_esw_port_status Some Port Lost Connection online_mask" between CBO and CBI, But the issued mask-bit was directed to an none-existed FEB. [PR1148869](#)
- When using MX Series platform as Layer 2 Tunnel Protocol (L2TP) L2TP access concentrator (LAC), if login/logout tunneled PPPoE subscribers over an extended period (e.g. login/logout 16K subscribers for 24- 48 hours), kernel crash may occur due to next-hop issue. [PR1150316](#)
- The outbound PPPoE control packets sourced from PPPoE daemon, such as PADO, error PADS and PADT, are assigned to queue 0 instead of queue 3. [PR1154070](#)

- Customer may see errors when doing 'show interface interface-set queue <if set>' for a pure numeric interface-set name. router> show interfaces interface-set queue 803 error: can't decode interface name `803': invalid device name. [PR1154667](#)
- Internal timing for bringing FPCs online is extended for MX 2020/2010 systems to accommodate longer initialization times for fabric and FPCs. [PR1164147](#)

Layer 2 Features

- In LDP Hierarchical VPLS (H-VPLS) topology (for example, the Multi-Tenant Unit switch (MTU-s) is connected to two PE devices via a primary spoke PW and backup spoke PW), when the primary spoke PW is down, an LDP address withdraw message with TLVs 0x404 and 0x405, which means "flush-all-from-me", will be sent from the PE (for example, PE1) on detection of failure of the primary spoke PW to peer PE devices participating in the full mesh to flush the MAC addresses learned in the corresponding Virtual Switch Instance (VSI). After receiving the message by a PE (for example, PE2) with "mac-flush propagate" configuration statement configured, the expectation is propagating "flush-all-from-me" to other participating PE (for example, PE3), but instead, it sends 'flush-all-but-me' message incorrectly. Because of this, the receiving PE (for example, PE3) will flush all MAC entries it learned, except the ones that were learned from LSI interface to sending PE (for example, PE2). [PR1131439](#)
- In VPLS scenario with AE interfaces as core facing interfaces, when LDP mesh-group is enabled with local-switching enabled in it, the neighbors configured under the local-switching hierarchical will cause LSI (Label-Switched Interface) to be created automatically. If port flapping occurs causing MPLS interface change associated with the LSI interface, the VPLS split-horizon might not be in functionality, this will cause traffic to be looped back. As a workaround, configuring configuration statement "enhanced-ip" can avoid this issue. [PR1138842](#)
- When configuring the "ecmp-alb" configuration statement to enable adaptive load balancing for equal-cost multipath (ECMP) next hops, the VPLS broadcast, unknown unicast, and multicast (BUM) traffic might be dropped on egress Packet Forwarding Engine when ingress/egress interfaces are distributed to more than one Packet Forwarding Engines. As a workaround, we can disable "ecmp-alb" to avoid this issue. [PR1142869](#)
- In a VPLS scenario, when "\$junos-underlying-interface-unit" is configured in "dynamic-profiles" hierarchy, which is then implemented in a routing-instance. The upgrade/commit will fail with the following error message, Parse of the dynamic profile (<dynamic_profile_name>) for the interface: \$junos-interface-ifd-name and unit: \$junos-underlying-interface-unit failed! [PR1147990](#)
- For routers equipped with the following line cards: T4000-FPC5-3D MX-MPC3E-3D MPC4E-3D-32XGE-SFPP MPC4E-3D-2CGE-8XGE MPC5E-40G10G MPC5EQ-40G10G MPC6E MX2K-MPC6E. If the router is working as VPLS PE, due to MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. [PR1148971](#)
- In subscriber management environment, when login/logout the subscribers, if the accounting feature is enabled as well as the underlying interface is configured with dynamic VLAN (DVLAN), the memory leak in "/mfs" may occur due to incorrect

interaction between Packet Forwarding Engine process (pfed) and authentication process (authd). [PR1112333](#)

- There is a bug in code of handling the redistribution of PPM (periodic packet management) Transmit and Adjacency entries for LACP, when the Interface entry is in pending distribution state. This issue might cause pppmd crash after graceful Routing Engine switchover. [PR1116741](#)
- For Routing Engine generated packet with VLAN tag, if the outgoing interface is an LT interface, the VLAN tag will not be removed even the LT interface is configured with untagged encapsulation. [PR1118540](#)
- For PVSTP/VSTP protocols, when MX/EX92xx router inter-operates with Cisco devices, due to the incompatible BPDU format (there are additional 8 Bytes after the required PVID TLV in the BPDU for Cisco device), the MX might drop these BPDUs. [PR1120688](#)
- In the DHCPv4 or DHCPv6 relay environment with large scaled environment (in this case, 50-60K subscribers), and the system is under stress (many simultaneous operations). The subscribers might get stuck in RELEASE state with large negative lease time. [PR1125189](#)
- In scenario that DHCP relay is used along with Virtual Extensible Local Area Network (VXLAN), if DHCP discover packet is received with the broadcast bit set via a VXLAN interface on MX platform (which is acting as DHCP relay), the OFFER back from the DHCP server will not be forwarded back to the client over the VXLAN interface. Unicast offers (that is, DHCP offer packet with unicast bit set) over VXLAN and both broadcast and unicast offers over native VLAN interfaces work fine. [PR1126909](#)
- In some rare scenarios, the MVRP PDU might unable to be transmitted, which could cause memory leak in layer 2 control plane daemon (l2cpd), and finally results in the l2cpd process crash. [PR1127146](#)
- Input/Output pps/bps statistics might not be zero after a member link of AE interface with distributed pppmd was down in M320/T-Series(GIMLET/STOLI based FPC). [PR1132562](#)
- The "Node ID" information is not shown on MX platform when traceoption flag "pdu" is configured to trace Ethernet ring protection switching (ERPS) PDU reception and transmission. [PR1157219](#)
- DHCP relay with forward-only cross-VRF results in bad packet format of the DHCP DISCOVER packet. Wireshark decode of packets from MX Series to DHCP server indicate Error; End options missing. [PR1157800](#)

MPLS

- With egress protection configured for Layer 3 VPN services to protect the services from egress PE node failure in a scenario where the CE site is multihomed with more than one PE router, when the egress-protection is un-configured, the egress-protection route cleanup is not handled properly and still point to the indirect composite nexthop in kernel, but the composite nexthop can be deleted in rpd even the egress protection route is pointing to the composite nexthop. This is resulting in composite nexthop "File exists" error when the egress protection is re-enabled and reuse the composite nexthop (new CNH addition fails as old CNH is still referenced in kernel). [PR954154](#)
- In MPLS scenarios, removing the "family mpls" configuration from an outgoing interface may cause inet and/or inet6 nexthops associated with that interface to unexpectedly transit to dead state. Even adding back "family mpls" cannot restore it. [PR1067915](#)
- Junk characters are being displayed in output of show connections extensive command. [PR1081678](#)
- When an LSP is link-protected and has no-local-reversion configured, if the primary link (link1) is down and LSP on bypass (link2), then another link (link3) is brought up, before the LSP switch to link3, if link1 is enabled and link3 is disabled, the LSP will stuck in bypass LSP forever. This is a timing issue. [PR1091774](#)
- If LDP is enabled via the 'protocols ldp' configuration option on a device running Junos OS, receipt of a spoofed, crafted LDP packet may cause the RPD routing process to crash and restart. [PR1096835](#)
- From Junos OS Release 13.2R1 and later, in MPLS L3VPN scenario, when the "l3vpn-composite-nexthop" configuration statement is enabled on a PE router and an interface style service set is attached to the ingress interface, the L3VPN packets with the MPLS labels will be sent to the service card and dropped. As a workaround, we should disable "l3vpn-composite-nexthop". [PR1109948](#)
- If "optimize-timer" is configured under P2MP branch LSP, this branch LSP will not be re-established if link flap on egress node. If "optimize-timer" is configured at protocols/mpls level, issue could be avoided. [PR1113634](#)
- For advertising IPV6 packets over MPLS GRE tunnel, the IPV6 address gets stuck in KRT queue. [PR1113967](#)
- For an MPLS L3VPN using LDP-signaled LSPs, in a rare racing condition (e.g. large-scale environment or Routing Engine CPU utilization is high), the rpd process might crash after an LDP neighbor down. [PR1115004](#)
- If an RSVP LSP has both primary and secondary standby path and link-protection enabled, a /32 bypass route is unhidden when the primary link goes down. This /32 route is supposed to be made hidden again when primary link comes back up. But in some cases, due to software defect, this /32 bypass route remains unhidden forever which causes some issues, for example, BFD session down due to better prefix received from Bypass LSP. [PR1115895](#)
- When multipoint LDP (M-LDP) in-band signaling is enabled to carry multicast traffic across an existing IP/MPLS backbone and routing process is enabled to use 64-bit mode, the rpd might crash due to accessing uninitialized local variables. [PR1118459](#)

- When OSPF LFA is enabled and there is available backup path, after clearing the LDP session to the primary path or backup path, in a very rare condition, the LDP session on this router might flap multiple times. [PR1119700](#)
- When local bandwidth accounting for inactive/adaptive standby path figures that there is not enough bandwidth (due to double-counting BW on common link shared by primary and secondary path) to fit it in an already full link and brings it down, CSPF will not be retried on the path unless there is some change in TE database. [PR1129602](#)
- When an PLR is a non-Juniper router, Juniper ingress node might stay on the bypass tunnel and ignore the CSPF result. [PR1138252](#)
- When a link fails on an RSVP LSP which has link-protection or node-link-protection configured, the PLR (point of local repair) will initiate a bypass LSP and the RSVP LSP will be tunneled on this bypass LSP. However, if now the bypass LSP is brought down because there is a link failure on it, the PLR might only send out session_preempted PathErr message to the upstream node without sending ResvTear message. Hence the ingress node does not receive ResvTear message and the RSVP LSP is not immediately torn down. The RSVP LSP will remain UP for more than 2 minutes until the RSB (Resv state block) on the ingress's downstream node gets time out and it sends ResvTear message to the ingress. [PR1140177](#)
- There is no entropy label for LDP route in scenario of LDP tunneling across a single hop RSVP LSP with label 0 (explicit-null) used. As workaround, either remove LDP tunneling or RSVP explicit-null will resolve the issue. [PR1142357](#)
- This issue is related to inter-op between multi vendor scenario. This fix will add sub-object RRO which will help change of label during FRR active scenario. [PR1145627](#)
- With NSR enabled and LDP configured, the rpd process may crash and restart on the new master Routing Engine after a Routing Engine switchover. [PR1155002](#)

Network Management and Monitoring

- On Junos OS Release 13.1X42/14.1X51/15.1R1/15.1R2, the SNMP average response time in the output of "show snmp statistics extensive" is wrongly calculated and might be observed with negative value. [PR1112521](#)
- While the router is rebooting and SNMP polling is not stopped, SNMP requests might land on mib2d process before Routing Engine protocol mastership is resolved, causing the mib2d process crash. [PR1114001](#)
- The SNMPv3 message header has a 4-byte msgID field, which should be in (0....2147483647), when the snmpd process has been running for a long time, the msgID might cross the RFC defined range and causing Net-SNMP errors, "Received bad msgID". [PR1123832](#)
- From Junos OS Release 14.1R1, SNMP informs are not sent out to the network management system (NMS) when significant events occur on a Junos device. As a workaround, we can configure an dummy trap-group. [PR1127734](#)

- A merge conflict was incorrectly resolved by changing snmp trap value of jnxDomLaneNotifications to 26. The correct value will always be 25. [PR1145144](#)
- With Junos OS Release 13.3R8/14.1R6/14.1X53-D30/14.2R5/15.1R2/15.1X49-D30 and later, when we configure fxp0 "master-only" address as source address of snmp trap, the snmp trap packets are not sent out after Routing Engine switchover. To restore this issue, we can use "restart snmp" or "delete/set snmp trap-options". As a workaround, we can use other addresses for snmp trap source. [PR1153722](#)

Platform and Infrastructure

- In the Network Time Protocol (NTP) configuration, if the specified source ip address is not in current routing-instance, the router will use primary address of interface (which will be used to send packet) as source address, Client routers will treat the NTP packets as incorrect packets, and then NTP synchronization failed. [PR872609](#)
- On MX Series based line cards, when GRE keepalive packets are received on a Packet Forwarding Engine that is different from the tunnel interface hosted, the keepalive message will apply the firewall filter configured on default instance loopback interface. [PR934654](#)
- Bad udp checksum for incoming DHCPv6 packets as shown in monitor traffic interface output. The UDP packet processing is normal, this is a monitor traffic issue as system decodes checksum=0000. [PR948058](#)
- When using MX2020 platform in Virtual Chassis (VC) environment, if the Virtual Chassis port (VCP) is located on the local Packet Forwarding Engine whose number is greater than 63 (i.e. VCP is located on local slot 16 or higher), the multicast traffic that should be sent to VCP will be dropped internally due to software issue. As a workaround, please configure the VCP ports on local chassis (local MX2020) slot 0 to 15, not 16 or higher. [PR1008676](#)
- When the 'enhanced-hash-key services-loadbalancing' feature is used by MX Series based line cards, load balancing of flows across multiple service PICs via the source-address across does not work when internal BGP (IBGP) is used to steer traffic to the inside service-interface. For example the operator will see on the stateful firewall that the same source-address has flows across multiple service interfaces. [PR1034770](#)
- When one of the "deny-commands" is incorrectly defined in the profile of TACACS+ server, all "deny-commands" regexes will be ignored, which leads to an over-permissive profile without any warning. [PR1078238](#)
- If with both MPC/MSDPC and other type of DPCs equipped, for local switching at mesh group level, split horizon on PW interfaces will not work and this would cause packets to loop back to same PW interface. [PR1084130](#)
- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)
- In 64-bit Junos OS environment, the Representational State Transfer (REST) API fails to start when configured with "set system services rest ...". [PR1097266](#)
- On MX Series platform, if ingress "multicast-replication" is configured, the throughput of the multicast may get reduced due to unnecessary threads during Packet Forwarding

Engine operation. In addition, only the performance of multicast traffic may get influenced (some of the multicast packets may get dropped on the Packet Forwarding Engine) by the issue. This PR has fixed/enhanced the performance. Now the performance limit should only be capped by fabric bandwidth in ingress Packet Forwarding Engine. In addition, before this fix, there was a limitation that VPLS/Bridging cannot run with ingress-replication feature as its BUM traffic cannot be handled by ingress-replication feature. This PR removed that limitation as well. Now BUM traffic for VPLS/Bridging is following normal multicast replication path even with ingress-replication feature. [PR1098489](#)

- With ECMP-FRR enabled, after rebooting the FPC which hoisting some ECMP links, the ECMP-FRR might not work. Clear any of BGP sessions (that is the part of ECMP) could help to clear this issue. [PR1101051](#)
- The kernel next-hop acknowledgement timeout maximum interval configured (krt-nexthop-ack-timeout) under the CLI hierarchy "routing-options forwarding-table" has been increase to 400 seconds to avoid performance issues with scaled subscribers. [PR1102346](#)
- On an MPC3E or MPC4E or on an EX9200-2C-8XS line card, when the flow-detection feature is enabled under the [edit system ddos-protection] hierarchy, if suspicious control flows are received, two issues might occur on the device: ? The suspicious control flow might not be detected on the MPC or line card. ? After suspicious control flows are detected, they might never time out, even if traffic flows no longer violate control parameters. [PR1102997](#)
- The following fields have been added to v10 Sampling (IPFIX) template and data packets: - SAMPLING RATE - SAMPLING INACTIVE TIMEOUT - SAMPLING ACTIVE TIMEOUT - TOTAL PACKETS EXPORTED - TOTAL FLOWS EXPORTED. [PR1103251](#)
- Improved VTY commands to show internal JNH memory usage. [PR1103660](#)
- On MX Series Virtual Chassis (MX-VC) with "locality-bias" configured, when equal-cost multipath (ECMP) load-balancing is occurring in the VC system, multicast streams and flooded Layer 2 streams may be duplicated or lost. As a workaround, we can disable "locality-bias" if possible. [PR1104096](#)
- Junos defines SNMP ifXTable (ifJnxInErrors/ifJnxInL3Incompletes) counter as 64-bit width, but it worked as 32-bit width counter. It works as 64-bit width counter after the fix. [PR1105266](#)
- Any configuration or logical interface (IFL) change will introduce 160 bits (20 bytes) memory leak on MPC heap memory when we have any type of inline sampling configured (ipfix or version 9). Only trigger of issue is the configuration of inline sampling, even without traffic being sampled. The leak is more evident in a subscriber management scenario when we have many IFL addition/deletion. Rebooting MPC in a controlled maintenance window is the only way to restore memory. [PR1105644](#)
- On MX Series-based platforms, in MX Series Virtual Chassis (MXVC) environment, if the subscriber logical interface (IFL) index 65793 is created (for example, when carrying 15K DHCPv4 subscribers to exceed IFL index creation 65793) and the IEEE 802.1p rewrite rule is configured (for example, using CoS rewrite rules for host outbound traffic), due to usage of incorrect IFL index, the Virtual Chassis Control Protocol Daemon

(vccpd) packets (for example, Hello packets) transmission may get lost on all VC interfaces, which may lead to VC decouple (split brain state, where the cluster breaks into separate parts). As a workaround, either delete the rewrite rule (delete class-of-service host-outbound-traffic ieee-802.1 rewrite-rules), or find the IFL in jnh packet trace that is not completing the vccpd send to other chassis and at Routing Engine clear that subscriber interface may resolve the issue. [PR1105929](#)

- When a common scheduler is shared by multiple scheduler maps which applies to different VLANs of an Aggregated Ethernet (AE) interface, if the configuration statement "member-link-scheduler" is configured at "scale", for some VLANs, the scheduler parameters are wrongly scaled among AE member links. As a workaround, we should explicitly configure different schedulers under the scheduler maps. [PR1107013](#)
- CVE-2015-5477 A vulnerability in ISC BIND's handling of queries for TKEY records may allow remote attackers to terminate the daemon process on an assertion failure. See this [KB article](#). [PR1108761](#)
- DHCP End options (option 255) is missing by DHCP-relay agent (where 20 bytes DHCP options 82 inserted) for client DHCP discover message with 19 bytes padding. [PR1110939](#)
- On MX-VC, when traffic with TPID 0x88a8 or 0x9100 is sending over AE interface, the packets which across VCP links might be dropped on egress VCP Packet Forwarding Engine due to invalid fabric token. [PR1112752](#)
- When inline BFD sessions and inline jflow are configured on the same Packet Forwarding Engine, with the increasing of active flows (about 65k), the BFD session might flap constantly and randomly due to the outgoing BFD packets are dropped. [PR1116886](#)
- Inline 6rd and 6to4 support for XL and XL-XM based platforms. [PR1116924](#)
- On MX Series-based FPC, when MPLS-labeled fragmented IPv6 packets arriving at PE router (usually seen in 6PE and 6VPE scenario), the Packet Forwarding Engine might mistakenly detect such IPv6 header and then drop these packets as "L3 incompletes" in the output of "show interface extensive". [PR1117064](#)
- When inline static NAT translation is used, if two rules defined in two service sets are pointing to the same source-prefix or destination-prefix, changing the prefix of one of the rule and then rolling back the changes is not changing back all the pools correctly. [PR1117197](#)
- On MX Series-based line cards, the firewall filter may have some issues when matching on Authentication Header (AH) protocol. This can affect VRRP (among others) when authentication is used, and an Routing Engine firewall filter is matching on protocol AH. As a workaround, we can change the filter to match on other criteria (e.g. source or destination address). [PR1118824](#)
- Tnetd is a daemon used for internal communication between different components like Routing Engine and Packet Forwarding Engines. It is used mainly to initialize the right server for rsh, rcp, rlogin, tftp, or bootp clients. It might crash occasionally due to the tnetd process not handling signals properly. [PR1119168](#)
- After changing an outer vlan-tags, the ifl is getting programmed with incorrect stp state (discarding), so the traffic is getting dropped. [PR1121564](#)

- With "fast-synchronize" configured, adding a new configuration-group that has configuration relevant to the rpd process and apply it and commit, then any configuration commits might cause the rpd process on the backup Routing Engine crash. We can reboot the backup Routing Engine to restore. [PR1122057](#)
- MX2020 or MX2010 running Junos OS software version 15.1 may experience "Minor" alarm associated with "i2c accelerator" timedout messages. [PR1122821](#)
- On MX Series-based platforms, when fragmented packets go through the inline NAT (including source NAT, destination NAT, and twice NAT), the TCP/UDP checksum would not be correctly updated. In this situation, checksum error would occur on the remote end (inside and outside device). Non-fragmented packets would not be affected by the issue. If possible, this issue could be avoided by either of the following workarounds, * Enable "ignore-TCP/UDP-Checksum errors" at the inside or outside device which processes TCP/UDP data OR * Make sure there will not be any fragments subjected to inline NAT functionality by appropriate MTU adjustment or setting. [PR1128671](#)
- Parity error at ucode location which has instruction init_xtxn_fields_drop_or_clip will lead to a LU Wedge. LU is lookup ASIC inside the MX Series. The LU wedge will cause the fabric self ping to fail which will lead to a FPC reset. This is a transient HW fault, which will be repaired after the FPC reset. There is no RMA needed unless the same location continues to fail multiple times. [PR1129500](#)
- On Junos OS devices with DHCP Relay config but without accounting config, and the accounting license does not exist, when the first DHCP control traffic is received, the following subscriber-accounting license grace period alarms might be triggered: alarmd[1650]: Alarm set: License color=YELLOW, class=CHASSIS, reason=License grace period for feature subscriber-accounting(30) is about to expire craftd[1592]: Minor alarm set, License grace period for feature subscriber-accounting(30) is about to expire. [PR1129552](#)
- For IPv6 packet with "no next header" in Hop-By-Hop header, if the Hop-By-Hop header length field value is large than 112, the router will drop such packet and log the following error: PPE PPE HW Fault Trap: Count 105, PC 60ce, 0x60ce: ipv6_input_finished_parsing LUCHIP(3) PPE_10 Errors lmem addr error. [PR1130735](#)
- NTP.org published a security advisory for thirteen vulnerabilities in NTP software on Oct 21st, 2015. These vulnerabilities may allow remote unauthenticated attackers to cause Denial(s) of Service(s), disruption of service(s) by modification of time stamps being issued by the NTP server from malicious NTP crafted packets, including maliciously crafted NTP authentication packets and disclosure of information. This can impact DNS services, as well as certificate chains, such as those used in SSL/https communications and allow attackers to maliciously inject invalid certificates as valid which clients would accept as valid. Refer to JSA10711 for more information. [PR1132181](#)
- Doing a file copy from a Routing-Engine running legacy Junos OS image to a Routing-Engine running Occam based Junos OS image fails. [PR1132682](#)
- Too many duplicate ACK messages are generated from Packet Forwarding Engine for TCP control connection with Routing Engine. This could cause: 1. MX-VC DDoS protection violation for VC-control low queue and makds MXVC split. 2. Cause Routing Engine and FPC high CPU utilization. [PR1133293](#)

- With scaled firewall filters attached to interfaces (e.g., 10k+ filters), running the "show configuration" command can cause high CPU of the mgd process. As a workaround, we can use the "show configuration |display set" command to view the configuration. [PR1134117](#)
- On XM chip based line cards (e.g. MPC3/4/5/6, and FPC type 5), in rare situation, when LU or XL chip congestion occurs (e.g. may occur when configuring with more than 4000 entries in the multicast list and large traffic performing replication, please note this is not a realistic configuration), XM chip wedge may occur. [PR1136973](#)
- On MX2020, when we remove whole power of a power zone, and then put the power back to the zone, FANTray LED stays Amber and FANTray LED on craft card stays OFF, and do not revert to green (FANTray LED) or ON (Craft LED) until we reboot the entire chassis system or hot swap that FAN tray. For Zone 0 (PSM 0 to 8), FAN 1 shows the above described behavior. For Zone 1 (PSM 9 to 17), FAN 3 shows the above described behavior. [PR1138209](#)
- On MX Series platforms with MX Series base line card, si interface is configured (i.e., set chassis fpc 1 pic 2 inline-services bandwidth 1g) and service is configured on the si interface. If Physical Interface is deleted while service is still configured, the FPC might crash. [PR1139348](#)
- On ungraceful exit of telnet (quit/shell logout), perm and env files created by pam were not deleted. [PR1142436](#)
- When the CLI command "show pfe statistics exceptions | match reject" is executed CPROD thread in the Packet Forwarding Engine may hog the CPU and result in FPC crash. [PR1142823](#)
- In certain affected Junos OS releases, executing "nhinfo -d" shell command might trigger a kernel panic. This is caused by insufficient buffer space in the routing socket requested by the "nhinfo" utility. [PR1148220](#)
- On MX2010 and MX2020 platforms, when error that causes adapter card (ADC)/Switch Fabric Board (SFB) initialization failure occurs (e.g., when Switch Processor Mezzanine Board (SPMB) is bringing up the ADC/SFB which has hardware issue), the SPMB crash occurs. [PR1149910](#)
- When the NTP server address is configured in Routing Instance table and reachable from inet.0 by static configuration (for example, by configuring static/route/next-table/VRF.inet.0), and NTP source-address is configured, the ntpd (the Network Time Protocol daemon running on NTP client) might pick the wrong source-address instead the configured source-address. As a result, NTP server cannot reply the NTP packet back. [PR1150005](#)
- Two interrupts are received from the FPGA on the control board of the MX2010/MX2020 platforms for every i2c transaction triggered from software. Only one is expected. [PR1151674](#)
- On MX Series routers with Junos OS Release 14.2R5-S1, when we specify a multiservice (ms-) interface to add a timestamp to Real-time Performance Monitor (RPM) probe messages, it will cause the mspmand process crash and the MS-MPC/MS-MIC keep crashing. As a workaround, we should configure RPM to perform timestamping either

on the Routing Engine (Routing Engine based RPM) or on an installed MPC Packet Forwarding Engine (Inline-RPM). [PR1152785](#)

- Fixed an issue with Inline Jflow where the Observation Domain field in exported IPFIX datagrams were always using the value attributed for LU0 in MPCs with multiple LUs per forwarding-engine. [PR1152854](#)
- Fixed an issue on where MX Series cards could crash while programming a firewall filter containing flexible-match-mask. [PR1157759](#)

Routing Protocols

- On large-scale BGP RIB, advertised-prefixes counter might show the wrong value due to a timing issue. [PR1084125](#)
- With this change the default label hold timer was increased for 10 seconds to 60 seconds. [PR1093638](#)
- When a BGP session supports multiple address families, the inactive route of some of the address families might not be flushed correctly, leading to wrong behaviors for some of the features which need to advertise inactive routes (e.g., advertise-inactive, advertise-external, optimal-route-reflection, etc). [PR1097297](#)
- Due to software bug, Junos OS cannot purge so called doppelganger LSP, if such LSP is received over newly formed adjacency shortly after receiving CSNP from the same neighbor. [PR1100756](#)
- IGMPv2 working in v2/v1 compatibility mode does not ignore v2 Leave messages received on a bridge-domain's L2 member interface. Moreover, an IGMP snooping membership entry for the respective group at this L2 member interface will be timed out immediately upon IGMPv2 Leave reception, even when there are some other active IGMP hosts attached to this L2 member interface. It might break multicast forwarding for this L2 member interface. [PR1112354](#)
- When two (or more) route target communities of MP-BGP route match to two (or more) route target communities in VRF import policy of a RI duplicate routing entries might be installed in the RI. In the output of 'show route table <RI name>.inet.0 detail' two identical routing entries appear with one being marked as 'Inactive reason: Not Best in its group - No difference'. When such duplicate routing information is to be deleted, rpd process process will crash. [PR1113319](#)
- When the Multicast Source Discovery Protocol (MSDP) is used, if the RP itself is the First-Hop Router (FHR) (i.e., source is local), the MSDP source active (SA) messages are not getting advertised by the RP to MSDP peers after reverse-path forwarding (RPF) change (e.g., the RPF interface is changed). [PR1115494](#)
- When a logical unit of an interface is associated with a Bidirectional Forwarding Detection (BFD) session, if changing the unit number of the interface (for example, change the unit number for a running BFD session from ge-1/0/0.2071 to ge-1/0/0.285), the device may fail to change the name due to the missing check for logical interface (IFL) index change. [PR1118002](#)
- On dual Routing Engine platform with nonstop active routing (NSR) and authentication of the Bidirectional Forwarding Detection (BFD) session enabled, BFD process (bfdd)

memory leak may occur on the master Routing Engine and the process may crash periodically once it hits the memory limit (RLIMIT_DATA). The problem does not depend on the scale, but the leak will speed up with more BFD sessions (for instance 50 sessions). As a workaround, if possible, disabling BFD authentication will stop the leak.

[PR1127367](#)

- When protocol MSDP is configured and then deleted, the NSR sync status for MSDP might stuck in "NotStarted", and ISSU might fail on master Routing Engine with reason "CHASSISD_ISSU_ERROR: Daemon ISSU Abort -1(NSR sync not complete: MSDP)".

[PR1129003](#)

- In multicast environment with Protocol Independent Multicast sparse mode (PIM SM) used, if a upstream router of last-hop router receives the (S,G) SPT join while the shortest-path tree (SPT) is not yet established (only because multicast source is not reachable, a reachable route for SPT which is just not established yet will not cause this issue), when the multicast route get deleted on the router (e.g., receives the (S,G) prune from downstream PIM router), the router would incorrectly stop forwarding the multicast traffic even if rendezvous-point tree (RPT) path exists. [PR1130279](#)
- On dual Routing Engine platforms, due to software issue, OSPF (including both OSPFv2 and OSPFv3) "DoNotAge" bit (e.g. source of LSA has flood-reduction feature enabled) is not mirrored to backup routing protocol process (rpd). In this situation, after performing nonstop active routing (NSR) switchover, the LSA on new master rpd remains without "DoNotAge" bit set. Once the LSA reaches OSPF max age, the router will flood LSA purge hence route flapping might be seen on all routers under the OSPF topology. [PR1131075](#)
- In rare condition, mt tunnel interface flap cause backup Routing Engine core. The exact root cause is not known. While processing updates on the backup Routing Engine (received from master Routing Engine), accessing free pointer cause the core. [PR1135701](#)
- On dual Routing Engine (Routing Engine) platforms with Bidirectional Forwarding Detection (BFD) protocol enabled, after graceful Routing Engine switchover (GRES), the periodic packet management process (ppmd) might crash on backup Routing Engine due to a software defect. [PR1138582](#)
- RPD generates core files while processing PIM hellos. There is no known workaround for this problem. RPD core seems to happen sometimes when a *g and sg's vanishes mostly due to LHR becoming a Non-DR from a DR. [PR1140230](#)
- With NSR configured, when the BFD sessions are replicated on backup Routing Engine, the master won't send the source address, instead backup Routing Engine will query the kernel to get the source address. In rare cases, the query might fail, resulting in the source address as all zeros. Later, if a GRES switchover happens, new master will have this all zeros source address. When BFD packet with this source address is send out, the other end will drop the BFD session due to no matching session (source address). [PR1145612](#)
- In the BGP labeled unicast environment, the secondary route is configured with both add-path and advertise-external. If the best route and secondary route are changed in a routing table at the same time, add-path might miss to readvertise the changed route. The old route with the old label is still the last route advertised to one router

instead of updating the advertisement with the new route and new label. So the traffic forwarding might be affected. [PR1147126](#)

- This core is seen because of incorrect accounting of refcount associated with the memory block which composes the nhid (IRB nh). When the refcount prematurely reaches to 0 we released the memory block while it was still referenced from a route. We may see this issue when mcsnoopd becomes a slow consumer of rtsock events generated by rpd (nexthop events in the current case) and messages get delivered in a out-of-order sequence causing the refcount to be incorrectly decremented. In the testbed where the issue was reported, tracing was enabled for mcsnoopd (for logging all events) causing it to become a slow consumer. However, it may become slow also for other reasons such as processing very high rate of IGMP snooping reports/leaves which could potentially trigger this to issue. [PR1153932](#)
- Core seen when BMP station was passive, and the BMP Collector was terminated non-gracefully, and BMP station was not properly cleaned up. [PR1154017](#)

Routing Policy and Firewall Filters

- When a malformed prefix is used to test policy (command "test policy <policy name> <prefix>"), and the malformed prefix has a dot symbol in the mask field (e.g., x.x.x.x/24), the rpd process might crash. [PR1144161](#)
- From Junos OS Release 13.2R1, an attempt to commit a configuration with a dangling conditional policy referring a non-existent/inactive routing-instance will be permitted. If we have a conditional policy referring an active routing-instance, deleting/deactivating this routing-instance and then committing will cause the rpd process crash. As a workaround, we should always make sure that conditional policies are referring active routing-instances. [PR1144766](#)

Services Applications

- The LCP state for tunneled subscriber is incorrectly displayed as "OPENED" (which reflects the LCP state before tunneling) by CLI command "show interfaces pp0.<unit>" on the LAC. This issue will be fixed from 15.1R3. As a workaround, we can use "show ppp interface pp0.<unit>" command to determine the correct LCP state for the subscriber. [PR888478](#)
- When polling to jnxNatSrcNumPortInuse via SNMP MIB get, it might not be displayed correctly. [PR1100696](#)
- Junos OS Release 13.3 and later releases, when configuring a /31 subnet address under a nat pool, the adaptive services daemon (SPD) will continuously crash. [PR1103237](#)
- SIP one way audio calls when using X-Lite SIP Softphone, in case that SIP media is switched to another media gateway though a SIP RE-Invite message. [PR1112307](#)
- In CGNAT environment, when a service PIC is in heavy load continuously, there might be a threads yielding loop in CPUs, which will cause the CPU utilization high, and might cause one the CPUs to be reset. [PR1115277](#)
- In CGNAT scenario, when we establish simultaneous TCP connects, we need to install timers for each TCP connection/flow. Due to this bug, we ended up creating two timers for the forward and reverse flow separately. Ideally there needs to be only one timer

for both the forward and reverse flow. Whenever the session used to get deleted due to timer expiry, the PIC used to crash whenever the code tried to delete the same flow again. [PR1116800](#)

- The Point-to-Point Tunneling Protocol (PPTP) ALG is used for tunneling Point-to-Point Protocol (PPP) packets over an IP network. But if the router configures session-limit-per-prefix, the PPTP-ALG does not work. [PR1128484](#)
- In L2TP environment, the max pass-through (or transit) sessions is 8192, due to there will be a delay to remove the session when receiving PADT messages from client, if there are mess logout during a short time, the limit might be reached and the jl2tpd will crash. This issue will affect the L2TP subscribers who is trying to login, the existing subscribers will not be affected. [PR1132285](#)
- With the following steps: 1) Define a RADIUS access profile with RADIUS which non-reachable from router 2) Run test access profile command 3) Abort using ctrl-c the l2tpd process will crash in few seconds. The existing active destinations, tunnels and sessions data will be recovered after the l2tpd restart. [PR1155345](#)

Software Installation and Upgrade

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

Subscriber Management and Services

- When the MX Series router acting as the Policy and Charging Enforcement Function (PCEF) uses Gx-Plus to request service provisioning from the Policy Control and Charging Rules Function (PCRF), the authentication service process (authd) might crash during the subscribers logout. [PR1034287](#)
- In a subscriber management environment, after scaling subscribers login/logout multiple times, the MX Series routers may hang the subscriber in the terminated state and be stuck in the backup accounting queue. The reason is that, when the authentication daemon (authd) is trying to fetch data from the session database (SDB), an error (for example, session not found, or an SDB deadlock or during the SDB recovery period) may occur, and this error may cause the router to fail to notify the client daemon to clean up the service records. In this case, the subscribers may not able to send Acct-Stop messages to the RADIUS server and end up staying in a terminated state. [PR1041070](#)
- This issue was introduced as part of another fix. Please contact JTAC for the recommended release for your deployment. [PR1049955](#)
- In the PPP environment, when a subscriber is logged out, its IFL index is freed, but in rare conditions the session database (sdb) entry is not freed. When the IFL index is assigned to a new IFL, it is still mapped to an old sdb entry, so the jpppd process might crash because of mismatching. The issue is not really fixed, developer just adds some debug information. [PR1057610](#)
- When using Neighbor Discovery Router Advertisement (NDRA) and DHCPv6 prefix delegation over PPPoE in the subscriber access network, if a local pool is used to allocate the NDRA prefix, when the CPE send DHCPv6 solicit message with both Internet Assigned Numbers Authority (IANA) and Identity Association Prefix Delegation

(IAPD) options, the subscriber might get IPv6 prefix from the NDRA pool but not the delegated pool. As a workaround, the CPE should send DHCPv6 solicit message with only IAPD option. [PR1063889](#)

- On MX Series platforms, in subscriber management environment, when receiving Activate-Service Vendor Specific Attributes (VSA) or Deactivate-Service VSA (for example, included in CoA-Request) from RADIUS server, the strings are parsed and empty "()" are stripped off service names, also, any white spaces are removed. Due to this reason, the service accounting message (e.g. "Accounting-Request") sent by the router (to the RADIUS server) does not contain empty "()" even if the strings were received in this way. As a workaround, changing RADIUS server to accept the service accounting message string without the "()" or the white space if possible. [PR1066709](#)
- In subscriber management environment with Remote Authentication Dial In User Service (RADIUS) server configured, when performing scaling subscribers login/logout, the device may stuck in RADIUS communication. [PR1070468](#)
- In subscriber management environment, the PPP daemon (jpppd) might crash repeatedly due to a memory double-free issue. [PR1079511](#)
- Activating and Deactivating services in same CoA-Req packet might fail to be executed on BNG router. Please note this issue will not be seen if there is no SRL service activated/deactivated request in this CoA. [PR1088366](#)
- In subscriber management environment with three or more radius-servers connected to an MX Series router, when AAA sends a request to one radius-server, if that particular request and all retries timeout, AAA records the time. For next request, AAA incorrectly uses the recorded time and marks that radius-server down even before trying to send out the request. [PR1091157](#)
- Radius backup accounting queue is used to store radius records while the radius server is not alive. Draining this queue when the server is reachable again should not log any critical message as this is normal operation. [PR1097491](#)
- On MX Series platforms, when using RADIUS dynamic requests for subscriber access management, if the device detects that the CoA-Request it received is same with the one in processing progress, the router would send CoA-NAK packet back to the RADIUS server with incorrect code 122 (invalid request) incorrectly. In this case, the router should return VSA with value "100 In Progress". [PR1100550](#)
- FFP is a generic process that will be called during commit process, and FFP calls the PDB initialization as part of its process. On the PDB-unsupported platforms (MX Series, EX9200, M10i, M120, M320 is PDB-supported), when committing configuration, some error messages will be seen. [PR1103035](#)
- In subscriber management environment, on MX Series platforms, if the configuration statement "last-statistics-when-unavailable" is configured, after the unrecoverable error, libstats is expected to not sending stats anymore, however, it is not the case here, the device may still send service interim-accounting message in wrong time-intervals to the RADIUS server. [PR1105954](#)
- On MX Series platforms, when using the DHCPv6 prefix delegation over PPPoE, if the RADIUS allocates a DHCPv6 pool name during the authentication of subscribers and "on-demand-ip-address" feature is enabled in a dynamic-profile, the prefixes may not

be cleared by authentication process (authd) after disconnecting the subscribers.

[PR1108038](#)

- When PPPoE sessions with Extensible Subscriber Services Management Daemon (essmd) subscribers configured, after terminating some PPPoE sessions without essmd service and executing a routing-engine switch, some PPPoE sessions cannot be set up. After terminating all sessions, some sessions are stuck in Terminating. The logout is queued because a Change of Authorization (CoA) is in progress and never complete. [PR1111062](#)
- On MX Series platforms, in subscriber management environment, if the sequence of event happens as following: the authentication process (authd) sends dynamic-profile service acct-start request to the Radius server (this is the service activated at login), then the CoA (for example, is used to activate the ESSM service) arrives at authd before the acct-start response, so the authd starts processing the CoA before processing the acct-start response, then during the processing of the acct-start response, the CoA, now in process, is deleted leaving authd with no way to answer the CoA request. As a result, the Radius server times-out and eventually sends a Disconnect request to authd, authd will deactivate any active services and deletes all of the subscriber's service entries (since the ESSMD services are not in the 'Active' state, so they are only deleted), at this point, the business 'subscribers' (interfaces) are orphaned and 'stuck'. The issue may be avoided by delaying the CoA requests by enough time to allow the authd to receive the acct-start responses for login. [PR1112323](#)
- When multiple authentication or accounting Radius servers are configured and if one of the servers is down/not-reachable, the Access-Request messages will be queued to the next Radius server no matter its "max-outstanding-requests" is reached or not. In case that all the Radius servers reached its "max-outstanding-requests", the new requests should be queued to an internal queue but they are queued to the last Radius server. As a workaround, we can use only one Radius server or make sure all the Radius servers are reachable. [PR1122703](#)
- In subscriber management environment, the authentication process (authd) crash may occur. This issue is not reproduced yet, possibly, it might be seen when generating a CLI Change of Authorization (CoA) request (e.g., via CLI command "request network-access aaa subscriber add service-profile filter-service session-id 10"), then logging out the subscriber (the one with service just activated), if the management CLI session is closed before subscriber entry is reused, the crash may occur. [PR1127362](#)
- In subscriber management environment with AAA authentication, after a few rounds of login/logout, some dynamic PPPoE subscribers might stuck in configured (AuthClntLogoutRespWait) state. [PR1127823](#)
- On MX Series platforms, with "subscriber-management" enabled, the authd process might crash during subscribers concurrent login/logout. When authd process crash, the new subscribers might not login. But all connected subscribers remain connected. The authd process will restore in a short time, then new subscribers could login successfully. [PR1128622](#)
- For Point-to-Point Protocol over Ethernet over ATM (PPPoEoA) subscribers, during subscriber bringing down, the assigned IFL unit number is not correctly retrieved, so it

can cause premature unit number exhaustion and thus fails to resolve *&junos-interface-unit/ &junos-interface-name* variables. [PR1137723](#)

- When class attribute is changed for a subscriber via COA, existing subscriber services continue to use the class attribute value at the time when that service was created. Updated class attribute value will take effect for the subscriber and the services created there. When both service and class attributes are present in COA request, AUTHD first processes the service requests and then processes class attribute. Due to this, accounting starts for requested services does not contain updated class attribute. [PR1143083](#)
- In normal BRAS environment, if the radius queue is presently full, MX BRAS might stop send accounting messages and customer might see "Radius result is CLIENT_REQ_MAXED_OUT" in authd log messages. [PR1152052](#)

User Interface and Configuration

- Junoscript traceoptions are available. [PR1062421](#)
- When committing a configuration with very long as-path, in this case the as-path is almost 12000 characters long, the commitd process might crash. The commitd process restart results in a minimal impact of system. As a workaround, please configure as-path less than 4096 characters long. [PR1119529](#)
- While using wildcard with interface like "set groups <group name> interfaces <xe> unit <unit>", there is no "disable" option followed. [PR1137377](#)
- When there are two or more sessions accessing the router, and one of the session (for example, session 1) is executing commit check in configuration private mode, if another session (for example, session 2) is keep executing commit and-quit in configuration private mode, because the commit check is not keeping the lock on local Routing Engine for entire session, there is a chance that session 2 will hit a Database opening error. The detailed sequence events are as following: (1) Session 1: commit check is not keeping the lock on local Routing Engine for entire session, once commit check on local is success, while it asked for lock on other Routing Engine. (2) Session 2: mgd acquired db lock on local Routing Engine. (3) Session 1: once commit check is completed on remote Routing Engine, it does cleanup and deleted the juniper.data+ (created by Session 2). (4) Session 2: juniper.data+ is still in use at local Routing Engine for by daemons and daemons start complaining about it and emitted the messages as "Database open failed for file '/var/run/db/juniper.data+' ". [PR1141576](#)

VPNs

- In NG-MVPN network, if there is a device working as PE which uses PIM, GRES/NSR Routing Engine switchover might cause multicast traffic loss. [PR1086129](#)
- In scenario involving pseudowire redundancy where CE facing interface in the backup neighbor (can be non-standby, standby, hot-standby type), if the virtual circuit (VC) is not present for the CE facing interface, the CE facing interface may go up after committing an unrelated VC interface configuration (e.g., changing description of another VC interface) even though the local pseudowire status is in down state. [PR1101886](#)

- On dual Routing Engine platform with BGP L2VPN and NSR configured, there might be a chance that the block label allocation and deletion for L2VPN is out of order on backup Routing Engine as following: Master rpd follows the below sequence (which is the correct order): Add Prefix P1 of Label L1 Delete Prefix1 of Label L1 Add Prefix P2 of Label L1 However, on backup rpd, it goes like this: Add Prefix P1 of Label L1 Add Prefix P2 of Label L1 <===== Delete Prefix1 of Label L1 In this situation, backup rpd cannot allocate the label L1 for P2 since L1 is already in use for P1, so it crashes. This occurs in scaling environment (10k L2VPN) where the router has multiple BGP peers and different L2VPN routing-instances are deleted and added back. [PR1104723](#)
- In Global Table Multicast (GTM) scenario (instance-type mpls-internet-multicast), when the GTM instance and master instance are used, if the name of the GTM instance is changed, the routing protocol process (rpd) may crash due to the usage of the incorrect routing table handle. [PR1113461](#)
- In L2circuit environment, if one PE has pseudowire-status-tlv configured but remote has not, and at the same time, this PE does not support control-word but remote does, then it will not send changed local status code to remote PE, in a rare condition, after enable status-tlv support at remote end, the L2circuit might stuck in "RD" state on remote PE. [PR1125438](#)
- In next-generation multicast virtual private network (MVPN) scenario, the rpd process will crash on the PE router after receiving PIM join messages from local receivers if "nexthop-hold-time" is configured in this local VPN routing and forwarding (VRF). As a workaround, we can disable "nexthop-hold-time" to avoid this issue. [PR1131346](#)

Resolved Issues: 15.1R2

- [Class of Service \(CoS\) on page 290](#)
- [Forwarding and Sampling on page 290](#)
- [General Routing on page 293](#)
- [High Availability \(HA\) and Resiliency on page 297](#)
- [Interfaces and Chassis on page 297](#)
- [Layer 2 Features on page 301](#)
- [MPLS on page 302](#)
- [Network Management and Monitoring on page 302](#)
- [Platform and Infrastructure on page 303](#)
- [Routing Policy and Firewall Filters on page 307](#)
- [Routing Protocols on page 307](#)
- [Services Applications on page 309](#)
- [Software Installation and Upgrade on page 310](#)
- [Subscriber Access Management on page 310](#)
- [User Interface and Configuration on page 311](#)
- [VPNs on page 311](#)

Class of Service (CoS)

- For an ATM interface configured with hierarchical scheduling, when a traffic-control-profile attached at ifd (physical interface) level and another output traffic-control-profile at ifl (logical interface) level, flapping the interface might crash the FPC. [PR1000952](#)
- In SNMP environment, when performing multiple walks or parallel snmpget for same interface at the same time (for example, SNMP bulk get/walk, or SNMP polling from multiple devices) on CoS related MIBs (jnxCos table), if the interface state changes or the request times out when FPC is responding the request, memory leak of Class-of-Service process (cosd) about 160 bytes (up to 1500 bytes) may occur, which may cause cosd to crash eventually when limit is exceeded. [PR1058915](#)
- On MX Series platform, when aggregate Ethernet (AE) interface is in link aggregation group (LAG) Enhanced mode, after deactivating and then activating one child link of the LAG, the feature that runs on AE interface rather than on the child link (for example, IEEE-802.1ad rewrite rule) may fail to be executed. [PR1080448](#)
- After restarting chassisd or doing an in-service software upgrade from 13.2R8.2 to 13.3R7.3, results in the following messages seen in syslog:
cosd_remove_ae_ifl_from_snmp_db ae40.0 error 2 Messages appear to be harmless with no functionality impact. [PR1093090](#)
- On MX104 platform, when we configure rate-limit for the logical tunnel (lt-) interface, the commit will fail. As a workaround, we can use firewall filter with policer to achieve the same function. [PR1097078](#)
- On MX Series platforms, when class-of-service (CoS) adjustment control profiles and "overhead-accounting" are configured, if the ANCP adjust comes before the logical interface (logical interface) adding message and the logical interface is in "UP" state when added (for example, it may occur when carrying scaling subscribers, for instance, 8K subscribers). For some of the subscribers, the local shaping rate from dynamic profile for the subscriber logical interface may not be overridden by shaping-rate of ANCP. [PR1098006](#)
- When performing the Routing Engine switchover without GRES enabled, due to the fact that the Class-of-Service process (cosd) may fail to delete the traffic control profile state attached to logical interface (IFL) index, the traffic-control-profile may not get programmed after the logical interface index is reused by another interface. [PR1099618](#)

Forwarding and Sampling

- When there are no services configured, datapath-traced daemon is not running. In the PIC, the plugin continues to try for the connection and continuous connection failure logs are seen. [PR1003714](#)
- In IP security (IPsec) VPN environment, after performing the Routing Engine switchover, the traffic may fail to be forwarded due to the SAs may not be downloaded to the PIC, or due to some security associations (SAs) on the PIC may incorrectly hold references for old Security Policy Database (SPD) handles while SPD has deleted its entries in the Security Association Database (SAD). [PR1047827](#)

- On all Junos OS based platforms, there are two different types of memory blocks that might be leaked. The first issue is rpd-trace memory block leak. There is one block each for any trace files opened for rpd. They could be leaked for each time a configuration commit is done. Around 40 bytes are leaked per operation. The issue does not occur in Junos OS Release prior to 14.1. The second issue is rt_parse_memory block leak which could happen during the configuration of aggregate routes, configuration information might not be freed. Around 16384 bytes are leaked per operation. This issue is a day-1 issue. [PR1052614](#)
- When enabling pseudowire subscribers the "show subscribers extensive" command does not display CoS policies applied to the subscriber interface. This issue was fixed in 13.3R6, 14.1R5 and 14.2R3. [PR1060036](#)
- For MX Series Virtual Chassis (MX-VC) with scaled subscribers, for example, 100K DHCP/20K PPPoE subscribers. If the Virtual Chassis port (VCP) FPCs also house the uplink ports and the "indirect-next-hop-change-acknowledgements" and "krt-nexthop-ack-timeout" configuration statements are configured along with the protection mechanism, after the master Routing Engine in the Virtual Chassis master router (VC-Mm) is powered down, the traffic loss and subscriber loss might be observed due to the indirect next-hop change acknowledgement timeout. With this fix, the upper limit for "krt-nexthop-ack-timeout" is changed from 100 seconds to 250 seconds. [PR1062662](#)
- For MX-VC platform, performing unified ISSU in scaled subscribers environment might cause all VC members to get restarted unexpectedly. [PR1070542](#)
- After rebooting the BNG with scaled subscribers, a dynamic-profile add request might fail, causing bbe-smgd (subscriber management daemon) to crash, then some subscribers might fail to login. [PR1071850](#)
- Juniper Networks device is not sending an error code to the Open vSwitch Database (OVSDb) client when the commit fails. Now a graceful mechanism is introduced to handle netconf configuration errors. If a netconf commit fails, the transaction will be routed to a failed queue. The transaction remains in the failed queue, until the user takes action to explicitly clear the transaction from the failed queue using the CLI. New CLI commands to show and clear failed netconf transactions. `user@router> show ovssdb netconf transactions Txn ID Logical-switch Port VLAN ID 1 vlan100 user@router> clear ovssdb netconf transactions` [PR1072730](#)
- On MX Series-based platform, when the Layer 3 packets destined to an Integrated Routing and Bridging (IRB) interface and then hit the underlying Layer 2 logical interfaces (IFLs), due to the egress feature list of the Layer 2 logical interfaces may get skipped, the features under the family bridge (for example, the firewall filter) on the Layer 2 interfaces may not be executed. [PR1073365](#)
- The issue is seen while moving an interface from one mesh group to another. [PR1077432](#)
- In scaled subscriber management environment (for example, 3.2K PPPoE subscribers), after heavy login/logout, the session setup rate keeps decreasing and also PAP-NAK messages are sent with "unknown terminate code". This continues till Broadband Network Gateway (BNG) does not accept PPP sessions and all newly incoming sessions are stuck in PAP Authentication phase (No PAP ACK received). [PR1075338](#)

- The license-check process may consume more CPU utilization. This is due to a few features trying to register with the license-check daemon which license-check would not be able to handle properly and results in high CPU on Routing Engine. Optimization is done through this fix, to handle the situation gracefully so that high CPU will not occur. [PR1077976](#)
- From Junos 14.1R1, if the hidden configuration statement "layer-4 validity-check" is configured, the Layer4 hashing will be disabled for fragmented IP traffic. Due to a defect, the Multicast MAC rewrite is skipped in this case, the fragmented multicast packets will be sent with incorrect destination MAC. [PR1079219](#)
- The rpd process might crash on both master and backup Routing Engines when a routing instance is deleted from configuration, if the routing instance is cleaned up before the interface delete is received from device control daemon (dcd). This is a rare timing issue. [PR1083655](#)
- OTN based SNMP Traps such as jnxFruNotifOperStatus and jnxIfOtnNotificationOperStatus are raised by offline/online MIC although no OTN interface is provisioned. [PR1084602](#)
- Invalid Ethernet Synchronization (ESMC) frames may be transmitted by MX router when activating LAG and tag-protocol-id under interfaces. [PR1084606](#)
- On a device with lt and ams interfaces configured, walking ifOutOctets or other similar OID's may cause a "if_pfe_ams_ifdstat" message to print. This is a cosmetic debug-level entry, which was incorrectly set to critical-level. [PR1085926](#)
- In the specific configuration of a LT interface in a VPLS instance and the peer-unit of this LT interface configured with family inet6 using vrrp, the kernel may crash when the FPC is online. [PR1087379](#)
- On MX Series based line card, if a rlsq interface is receiving continuous fragmented traffic, doing rlsq switchovers couple of times might cause FPC to crash and reboot. [PR1088300](#)
- In rare cases, SSH or telnet traffic might hit incorrect filter related to SCU (Source Class Usage) due to the defect in kernel filter match. This issue comes when the filter has match condition on source class ID. [PR1089382](#)
- In rare cases, MX Series routers might crash while committing inline sampling related configuration for INET6 Family only. [PR1091435](#)
- In a fib-localization scenario, IPv4 addresses configured on service PICs (SP) will not appear on FIB-remote FPCs although all local (/32) addresses should, regardless of FIB localization role, install on all Packet Forwarding Engines. There is no workaround for this and it implies that traffic destined to this address will need to transit through FIB-local FPC. [PR1092627](#)
- There are entries for PEM in jnxFruEntry in VMX. It is not necessary and is cosmetic. [PR1094888](#)
- When route convergence occurred, the new gateway address is not updated correctly in inline-jflow route-record table (route-record table is used by sampling), and the sampling traffic forwarding might be affected, but normal routing would be not affected. [PR1097408](#)

- After upgrading to Junos OS Release 14.1R1 and higher, loopback ISO family address may be stuck in KRT queue. [PR1097778](#)
- When BGP multipath is enabled in a Virtual Routing and Forwarding (VRF), if "auto-export" and "rib-group" are configured to leak BGP routes from this Routing Instance table to another, for example, the default routing table, then traffic coming from the default routing instance might not be properly load balanced due to the multipath-route leaked into the default routing table is not the active route. This is a random issue. As a workaround, only use "auto-export" to exchange the routes among the routing tables. [PR1099496](#)

General Routing

- There is hardware design flaw with 2x10GE MIC and 4x10GE MIC today which introduces +/-6.2ppm frequency offset for SyncE operation. In order to correct this, the framing of the PIC and interface has to be matched (which will not be by default). [PR932659](#)
- SNMP MIB walk of object "jnxSpSvcSet" gives hardcoded value as "EXT-PKG" for SvcType. [PR1017017](#)
- With Multiservices MPCs (MS-MPCs) or Multiservices MICs (MS-MICs) installed on MX Series platform, when trying to view the Network Address Translation (NAT) mappings for address pooling paired (APP) and/or Endpoint Independent Mapping (EIM) from a particular private or a public IP address, all the mappings will be displayed. [PR1019739](#)
- On MX Series router with MPC3E/MPC4E/MPC5E/MPC6E if the Packet Forwarding Engine has inline NAT configured or is processing inline GRE decapsulation with packet-sizes between 100B-150B, in some very corner cases, traffic blackhole might be seen due to incorrect cell packing handling. On T4000 with FPC type 5, when these cards are processing any packets sizes between 133B-148B in certain sequences causes incorrect cell packing handling. [PR1042742](#)
- Queue stats on LSQ interfaces are not properly cleaned up when queuing enabled on the Physical Interface and the queues hosted at Physical Interface level. This happens when a subsequent delete and create of LSQ interface (not always though) - 14.1R4.10. [PR1044340](#)
- MPC with Channelized OC3/STM1 (Multi-Rate) Circuit Emulation MIC (MIC-3D-4COC3-1COC12-CE) might crash. This problem is very difficult to replicate and a preventive fix will be implemented to avoid the crash. [PR1050007](#).
- In subscriber management environment, the Berkeley Database (DB) may get into deadlock state. It is brought on by multiple processes attempting to simultaneously access or update the same subscriber or service record. In this case, due to the access to DB were blocked by device control daemon (dcd), the subscriber management infrastructure daemon (smid) fails to recover the DB. Consequently, the router may stop responding to all the login/logout request as well as statistics activity. This timing related issue is most likely to occur during login or logout and when the system is busy. [PR1054292](#)
- With inline L2TP IP reassembly feature configured, the MX Series routers with MPCs/MICs might crash due to a memory allocation issue. [PR1061929](#)

- In subscriber management environment, if IPv6 family is not enabled in the dynamic profile, the IPv6 Router Advertisement message will not be sent through the dynamic subscriber interface. As a workaround, you can enable family inet6 in the dynamic profile. [PR1065662](#)
- When setting the syslog to debug level (any any), you may note reoccurring messages of the form "ifa for this rt ia is not present, consider ifa as ready". These messages are logged for IPv6 enabled interfaces when receiving forwarded packets and cause no harm. Set a higher debug level to avoid seeing them. [PR1067484](#)
- The static route prefers the directly connected subnet route for resolving the nexthop rather than performing a longest prefix match with any other available routes. In case of longest prefix route being desired in customer deployment, it will result in traffic loss issue. Now a new configuration statement "longest-match" is introduced to enable longest prefix matching behavior when desired: set routing-options static route <destination prefix> next-hop <address> resolve longest-match. [PR1068112](#)
- In subscriber management environment, changing the system time to the past (for example, over one day) may cause the processes (for example, pppoe, and autoconfd) that use the time to become unresponsive. [PR1070939](#)
- Higher baseline CPU utilization and periodic CPU spikes might be seen on XM-based MPC as compared to MPC-3D-16XGE-SFPP Cards due to the following reasons: On MX-based MPC, low priority threads which monitor various things in the background on a periodic basis such as voltage, temperature, stats counters, hardware status exist. When the system is idle, these threads are allowed to take more of the load and that is why higher baseline CPU/CPU spikes are seen. This does not prevent other higher priority threads from running when they have to, as these are non-critical activities being done in the background and hence it is a non impacting issue. [PR1071408](#)
- Traffic throughput test between MPC1/1E/2/2E card and MPC2E/3E NG card, the flowing from MPC1/1E/2/2E card to MPC2E/3E NG card is lesser then from MPC2E/3E NG card to MPC1/1E/2/2E card. [PR1076009](#)
- Vendor provided the fix, which includes conditional check. [PR1076369](#)
- In a Q-in-Q setup, if outer vlan tag is coming with EtherType 0x88a8, it is not possible to create dynamic vlan interface on Junos 13.1X42 or 14.1X51 releases. [PR1080734](#)
- On MX Series platform with MS-MPC/MS-MIC, in some mspmand process crash scenarios, after the mspmand coredump is finished or almost finished, PIC kernel also crashes and dumps vmcore. The mspmand generates core files in these scenario are readable but vmcores are not. [PR1081265](#)
- In DHCPv6 prefix delegation over PPPoE scenario, when forwarding the control packet from the Routing Engine to the DHCPv6 identity association for prefix delegation (IA_PD) address over PPPoE, for instance, executing ping from Routing Engine targeting the client's PD address, the traffic may get dropped on the device. [PR1081579](#).
- If a router has Service PIC equipped but without any Service PIC specific configurations, the CPU usage on this PIC/FPC might be high. Have some configurations under below configuration statement could prevent from this issue: [system processes process-monitor traceoptions] OR [chassis fpc <fpc slot> pic <pic slot> adaptive-services service-package extension-provider] OR [services] [PR1081736](#)

- In multi-homing and signal active EVPN scenario, if IRB interface is included in the instance, when the DF-CE link flaps, due to a timing issue, the DF might send L3 EVPN routes with label 0 to remote PEs, causing traffic to be dropped at remote PE. [PR1082287](#)
- 'show interfaces queue <ifl>' stats are not correct with RLSQ warm-standby mode. Issue seen on MPCs and MICs as well in 14.1R4.10. [PR1082417](#)
- TCP messages do not have their MSS adjusted by the Multiservices MIC and MPC if they do not belong to an established session. [PR1084653](#)
- With a scaled subscribers system, repeatedly doing tcpdump of subscriber interface and press ctrl+c might cause bbe-smgd daemon memory growing, which will in turn causing crash, SDB corruption and some other daemons crashing. Following signs may be seen when this problem is hit: log messages like: "/kernel: cmd bbe-smgd pid 1997 tried to use non-present sched_yield" tcpdump stops working bbe-smgd no longer accepts new vty sessions. [PR1085944](#)
- In some rare conditions, depending on the order in which configuration steps were performed or the order in which hardware modules were inserted or activated, if PTP master and PTP slave are configured on different MPCs on MX Series router acting as BC, it might happen that clock is not properly propagated between MPCs. This PR fixes this issue. [PR1085994](#)
- MACsec using static secure association key (SAK) security mode does not work properly on MX80 routers and FPC slots other than slot 0 of MX104 routers. [PR1086117](#)
- mspmand.core is observed while making ms-mic offline with IPsec and Jflow configured on same ms-mic with dynamic IPSEC tunnels. [PR1086819](#)
- If the ALG is receiving UDP fragmented control traffic (e.g. SIP control packets) continuously, the mspmand process (which manages the service PIC) might crash due to buffer error. [PR1087012](#)
- Due to a change in an existing PR, group names in the configuration must be a string of alphanumericals, dashes or underscores. There is not workaround other than following the group name instructions. [PR1087051](#)
- On LAC (L2TP Access Concentrator) router with session client-idle-timeout configured, the tunneled PPP session will always keep active due to the PPP control messages are accounting as user data. [PR1088062](#)
- Wrong ESH checksum computation with non-zero Ethernet Padding in Juniper MX Series router. [PR1091396](#)
- The mspmand process might crash due to prolonged flow-control with TCP ALGs under the following possible scenario, mostly when the following conditions happen together: 1. When the system is overloaded with TCP ALG Traffic 2. There are lots of retransmissions and reordered packets. [PR1092655](#)
- When the control path is busy/stuck for service PIC, the AMS member interface hoisted by it might be down, but when the busy/stuck condition is cleared, the member interface might not recover, and AMS bundle still shows the PIC as inactive. [PR1093460](#)

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

- On MX dual Routing Engine platforms, if there are a large number of addresses (in this case, there are > 500 addresses configured, the issue might be observed around 472 addresses) configured on lo0.0, when the Broadband Edge subscriber management daemon (bbe-smgd) replicating these addresses to the standby Routing Engine, the internal 8K replication buffer may get exceeded. Due to this failure, memory leak (around 45MB every time error is encountered) may occur when bbe-smgd tries to delete the object. Since lo0.0 object gets created/destroyed over and over, bbe-smgd runs out of memory and crash eventually. [PR1101652](#)
- Non-queuing MPC5E might crash continuously if rate-limit under transmit-rate for scheduler is applied. As a workaround, do not configure rate-limit and use firewall policer for forwarding-class instead. MPC5EQ is not exposed. [PR1104495](#)
- On MX Series platform, the output of CLI command "show system subscriber-management route" may be shown as empty. [PR1104808](#)
- An IPv4 filter configured to use the filter block with term that has both "from precedence" and another non 5-tuple (i.e. not port, protocol, address) will cause an XL/EA based board to reboot. Example: set firewall family inet filter FILTER fast-filter-lookup set firewall family inet filter FILTER term TERM from precedence PRECEDENCE set firewall family inet filter FILTER term TERM from tcp-established. [PR1112047](#)

High Availability (HA) and Resiliency

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

Interfaces and Chassis

- Chap Local-name default to 8 characters. Should be 32. [PR996760](#).
- On MX Series platform with large-scale PPPoE subscribers (more than 60k) connected, PPP client process (jpppd) might crash and generate core files when performing RE switchover. [PR 1018313](#)
- If a subscribers-facing AE interface has link protection enabled, offline the primary child link hosted FPC might cause some subscribers to down. [PR1050565](#)
- dcd will crash if targeted-distribution applied to ge ifd via dynamic-profile. [PR1054145](#)
- There is a mismatch in mac statistics, few frames go unaccounted. This is a day-1 issue with the software fetching of mac statistics, the snap and clear bits were setting set together on pm3393 chip driver software, so it used to so happen that even before the copy of stats to shadow registers happened, clear was happening which used to go unaccounted. [PR1056232](#)
- When a dynamic PPPoE subscriber with targeted-distribution configured on a dynamic vlan demux interface over aggregated ethernet, the device control daemon (dcd)

process might crash during a commit if the vlan demux has mistakenly been removed. The end users cannot visit internet after the crash. This is a rare issue and not easy to be reproduced. [PR1056675](#)

- It is observed that the syslog messages related to kernel and Packet Forwarding Engine may get generated at an excessive rate, especially in subscriber management environment. Most of these messages may appear repeatedly, for example, more than 1.5 million messages may get recorded in 2 hours, and there are only 140 unique messages. Besides, these messages are worthless during normal operation and due to the excessive rate of log generation, it results in high Routing Engine CPU consumption (for example, Routing Engine CPU utilization can be stuck at 100% for a long time (minutes or hours), it depends on the activity of subscribers (frequency of logins and logouts) and on the AI scripts used by the customer) by event process (eventd) might be observed on the device. [PR1056680](#)
- When "set chassis lcc 0 offline" is used on SCC and committed, the configuration gets synced on LCC. However when "delete chassis lcc 0 offline" is used on SCC, we need to do commit two times on SCC in order to sync the config on LCC being brought online. [PR1058994](#)
- For transit traffic on INLINE LSQ redundancy (rlsq) interface, the input firewall-filter counters are logging zero packet count regardless of traffic flow. Output filter counters are logging correctly. For host-bound traffic, the firewall output counter will get double accounted on Classical rlsq and triple accounted on INLINE rlsq. This issue is targeted to be fixed in Junos 14.1R5. [PR1060659](#)
- In scaling PPP subscriber environment, when the device is under a high load condition (for example, high CPU utilization with 90% and above), the long delay in session timeout may occur. In this situation, the device may fail to terminate the subscriber session (PPP or PPPoE) immediately after three Link Control Protocol (LCP) keepalive packets are missed. As a result, the subscriber fails in reconnect due to old PPP session and corresponding Access-Internal route are still active for some time. In addition to this, it is observed that the server is still sending KA packets after the session has timed out. [PR1060704](#)
- For Junos OS Release 13.3R1 or above, after multiple (e.g. 26) iterations of graceful Routing Engine switchover (GRES), the TNP address of management interface might be deleted incorrectly during switchover, this leads to all FPCs being offline. [PR1060764](#)
- Link Up/Down SNMP traps for AE member links might not be generated, but the SNMP traps for the AE bundle works well. [PR1067011](#)
- In PPP-based subscriber management environment, after performing scaling subscribers login/logout, the subscribers might be stuck in terminating and terminated state because logout requests are not processed properly, and the Session Database (SDB) might get exhausted eventually after stuck subscribers exceeding 256000. [PR1073146](#)
- In PPP subscriber management environment, the jpppd process might crash for a timing issue. [PR1074545](#)
- When the Ethernet Link Fault Management (LFM) action profile is configured, if there are some errors (refer to the configuration, for example, frame errors or symbol errors)

happening in the past (even a long past), due to the improper handling of error stats fetching from kernel, the LFM process (lfmd) may generate false event PDUs and send false alarm to the peer device. [PR107778](#)

- On MX Series Virtual Chassis (MX-VC) platform, due to a timing issue, the physical interface (ifd) on the same Modular Interface Card (MIC) with Virtual Chassis port (VCP) might not be created or takes a very long time to be created after rebooting the hosted Modular Port Concentrator (MPC). [PR1080032](#)
- MAX-ACCESS value has been changed in jnx-otn.mib for the following oids:
jnxOtnIntervalOdu15minIntervalNumber jnxOtnIntervalOtu15minIntervalNumber
jnxOtnIntervalOtuFec15minIntervalNumber The value has been changed from read-only to not-accessible to be inline with newer MIBs. [PR1080802](#)
- On MX Series platform acting as broadband network gateway (BNG), in Point-to-Point Protocol (PPP) scenario, when using the Internet Protocol version 6 Control Protocol (IPv6CP) for negotiation, if the router receives an IPv6CP Configure-Request packet from client, MX BNG sends the Configure-Request packet, but does not send IPv6CP Configure-Ack packet, in case it does not receive the Configure-Ack that responding to the Configure-Request packet it sent. The behavior does not follow the RFC 1661, which demands both the actions Send-Configure-Request (i.e. IPv6CP-ConfReq from MX to client) and Send-Configure-Ack (i.e. IPv6CP-ConfAck from MX to client) to be conducted on the router without any significant delay. [PR1081636](#)
- With Non-MX Series/service DPCs which are not supported with enhanced-ip, when these unsupported DPCs are in the chassis, the user switches to enhanced-ip and reboots the router, the router should come back up and the unsupported DPCs should stay powered off and not log any alarms. In this case, the non-supported DPCs stay powered off, but they are also continuing to raise alarms. There are two workarounds for this issue; first, power down the FPC prior to changing enhanced-ip mode; second, perform a hard restart by "restart chassis-control immediately" to restore. Both of these workarounds will impact traffic through the router. [PR1082851](#)
- In MX virtual chassis (MXVC) scenario, during unified ISSU operation, the new master Routing Engine does not have the MXVC SCC's system MAC address. It just has its local system MAC address. The address is not replicated between local Routing Engines, and the new master Routing Engine is not yet connected to the MXVC SCC to receive it. Hence, the possibility of overwriting the FPC with an address that does not match the previous address exists. [PR1084561](#)
- The VRRP preempt hold time is not being honored during NTP time sync and system time is changed. [PR1086230](#)
- On MX Series Virtual Chassis (MX-VC) platform with "subscriber-management" enabled, after power up/reboot, the VC backup router (VC-B) experiences a rapid sequence of role transitions from no-role to VC master router (VC-M) to VC-B, the expected local GRES and a reboot of the former master Routing Engine might not happen on the VC-B. Some of the FPCs on it might be stuck in "present" state and eventually rebooted. [PR1086316](#)
- Deactivating/activating logical interfaces may cause BGP session flapping when BGP is using VRRP VIP as the source address. This is caused by a timing issue between dcd

and VRRP overlay file. When dcd reads the overlay file, it is not the updated one or yet to be updated. This results in error and dcd stops parsing VRRP overlay file. [PR1089576](#)

- In the dual Routing Engines scenario with GRES and ae0 interfaces configuration, if GRES is disabled on system, the backup Routing Engine should remove the ae0 bundle, however it does not go clean and ae0 remains in backup Routing Engine. After switching Routing Engine mastership to make other Routing Engine as master, the new master Routing Engine (which was backup earlier) continues to use invalid MAC address "00:00:00:00:00:00". [PR1089946](#)
- When an interface on SFPP module in MIC is set disabled, after pulling out the SFPP and then insert it, the remote direct connected interface might get up unexpectedly. [PR1090285](#).
- After removing a child link from AE bundle, in the output of "show interface <AE> detail", the packets count on the remaining child link spikes, then if add back the previous child link, the count recover to normal. [PR1091425](#)
- For Junos OS version 14.1X51-D60 or 14.1X50-D105, when DHCP local server is configured, the DHCP subscribers might be unable to come up. [PR1092553](#)
- In MX Series Virtual Chassis (MXVC) environment, when rebooting the system or the line cards which contain all the Virtual Chassis port (VCP) links, because line cards might fail to complete the rebooting process within 5 minutes, the timer (that is, the amount of time allowed for the LCC to connect to the SCC) started by the master router might expire which may cause the VCP links establishment failure. In addition, this issue is not specific to the line cards type, based on the observation, the timer (5 minutes) may expire on a MX2020 with all 20 FPCs equipped as well. [PR1095563](#)
- On PB-2OC12-ATM2-SMIR PIC, port 0 and port 1 are configured with clock source as external, if Loss of signal (LOS) is inserted on port 0, the port 0 will go down, the expected behavior is clock being used from port 1. But in this case, port 0 down will results in port 1 flapping and reporting SONET phase lock loop (PLL) errors. [PR1098540](#)
- In VRRP environment, with VRRP configured over double tagged interface and VRRP delegate-processing enabled, the PDUs are generated with only one tag and the outer tag is not added, because of which, the PDUs will get dropped at the receiving end. The similar configuration that may cause the issue might be seen as below, .. protocols { vrrp { delegate-processing; <<<<< "delegate-processing" is enabled for VRRP } ... interfaces { xe-0/0/3 { flexible-vlan-tagging; unit 0 { vlan-tags outer 2000 inner 200; <<<<< VRRP is configured over double tagged interface family inet { address 10.10.10.147/29 { vrrp-group 17 { virtual-address 10.10.10.145; priority 100; accept-data; } } } } } .. [PR1100383](#)
- After configuring related ae interface configuration, we might find some of ae interfaces disappear in MX-VC. It seemed that ae interfaces are not allocated MAC address from chassisd properly. * This issue only happens in the first configuration timing after rebooting/restarting chassisd. So even if you configure related ae interface configuration repeatedly, you cannot find this issue. When this issue happens these message will be seen in the messages logs. -----
lab@router_re0> show log messages| match CHASSISD_MAC_ADDRESS_AE_ERROR
Jun 26 16:04:34.064 router_re0 scchassisd[2008]:
CHASSISD_MAC_ADDRESS_AE_ERROR: chassisd MAC address allocation error for

ae4 Jun 26 16:04:34.105 router_re0 /kernel: Jun 26 16:04:34.064 router_re0
 scchassisd[2008]: CHASSISD_MAC_ADDRESS_AE_ERROR: chassisd MAC address
 allocation error for ae4 ----- Restore ae
 interfaces * This is not workaround. deactivate/activate ae interfaces. (We need to do
 this to all disappeared ae interfaces.) [PR1100731](#)

- To ensure that the router or switch is reachable for management purposes while it boots or if the routing protocol process fails to start properly, we can configure a backup router, which is a router that is directly connected to the local router or switch (that is, on the same subnet) through its private management interface (for example, fxp0 or me0). When a backup router running IPv6 and a static route to reach the management network are configured, some invalid IPv6 routes are added to default forwarding-table on the master or the backup Routing Engine. [PR1100981](#)
- Due to the fact that the error injection rate configured by user on Routing Engine via CLI command "bert-error-rate" may not be programmed in the hardware register, the PE-4CHOC3-CE-SFP, PB-4CHOC3-CE-SFP, MIC-3D-4COC3-1COC12-CE, and MIC-4COC3-1COC12-CE-H may fail to inject bit errors during a Bit Error Ratio Test (BERT). [PR1102630](#)
- On MX240 or MX480 platform with at least two DC modules (PN: 740-027736) equipped, when shutting down one of the PEMs and then turn it on again, even the PEM is functioning, the "PEM Fan Fail" alarm might be observed on the device due to software logic bug. There is no way to clear the ALARM_REASON_PS_FAN_FAIL for I2C_ID_ENH_CALYPSO_DC_PEM once it has been raised. [PR1106998](#)

Layer 2 Features

- Under rare circumstances it is possible for the DHCP drop counts for reason SEND ERROR to be incremented twice for a single failure. [PR1009296](#)
- MTU change is not advised on the Ethernet ring protection (ERP) ring interfaces unless ring is in idle condition. Changing ring interface MTU while ring is not in idle state might result in change in the forwarding state of the interface which can lead to loop in the ring. [PR1083889](#)
- When family bridge was configured and committed, l2ald repeated restarting with core. After l2ald repeated restarting several times, it stopped working due to thrashing condition. Core of l2ald will be seen with the following configuration. set interfaces fxp0 unit 0 family bridge interface-mode access set interfaces fxp0 unit 0 family bridge vlan-id 100 When the configuration is committed, message like following is logged and core is generated. l2ald[1624]:

```

../././././src/junos/usr.sbin/l2ald/l2ald_vpls_flood.c:3117: insist '!err' failed l2ald[1734]:
../././././src/junos/usr.sbin/l2ald/l2ald_vpls_flood.c:3117: insist '!err' failed l2ald[1769]:
../././././src/junos/usr.sbin/l2ald/l2ald_vpls_flood.c:3117: insist '!err' failed l2ald[1993]:
../././././src/junos/usr.sbin/l2ald/l2ald_vpls_flood.c:3117: insist '!err' failed l2ald[2195]:
../././././src/junos/usr.sbin/l2ald/l2ald_vpls_flood.c:3117: insist '!err' failed ... init:
l2-learning is thrashing, not restarted PR1089358

```
- During interface flaps, a high amount of TCN (Topology Change Notification) might get propagated causing other switches to get behind due to high amount of TCN flooding. This problem is visible after the change done from Junos OS Release 11.4R8

and later, which propagates TCN BPDU immediately and not in the pace of the 2 second BPDU. Hello interval to speed up topology change propagation. The root cause is that the TCNWHILE timer of 4 seconds is always reset upon receiving TCN notifications causing the high churn TCN propagation. [PR1089580](#)

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

MPLS

- In Resource Reservation Protocol (RSVP) environment, if CoS-Based Forwarding (CBF) for per LSP (that filter out traffic not related to that LSP) is configured, and either the feature fast-reroute or link-protection is used on the device, when the primary link is down (for example, turning off the laser of the link), due to some next hops of the traffic may be deleted or reassigned to different class of traffic, and the RSVP local repair might fail to process more than 200 LSPs at one time, the traffic might get dropped by the filter on the device before the new next hop is installed. In this situation, the feature (fast reroute or link protection) may take longer time (for example, 1.5 seconds) to function and the traffic loss might be seen at the meantime. In addition, the issue may not be seen if the CBF for per LSP is not configured on the device. [PR1048109](#)
- In race conditions, the rpd process on backup Routing Engine might crash when BGP routes are exported into LDP by egress-policy and configuration changes during the rpd process synchronizing the state to backup rpd process. [PR1077804](#)
- On dual Routing Engine platform with GRES, the kernel synchronization process (ksyncd) might crash on the backup Routing Engine when adding of route pointing to indirect nexthop on system. [PR1102724](#)

Network Management and Monitoring

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

Platform and Infrastructure

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

- In Junos release 13.3R6 or 14.2R3, for PPPoE subscribers over the aggregated Ethernet (ae) interface, the output of "show interface statistics <pp> detail" command shows the ingress/egress traffic statistics for the aggregate interface instead of the statistics for PP/DEMUX logical interface. [PR1069242](#)
- Having "shared-bandwidth-policer" on an aggregated ethernet interface; if a member interface flapped, the NPC which the interface belongs may restart. Similar issue may also happen when changing the firewall policer configuration. [PR1069763](#)
- When Integrated routing and bridging (IRB) interface is configured with Virtual Router Redundancy Protocol (VRRP) in Layer 2 VPLS/bridge-domain, in corner cases after interface flapping, MAC filter ff:ff:ff:ff:ff:ff is cleared from the Packet Forwarding Engine hardware MAC table, so the IRB interface may drop all packets with destinations MAC address FFFF:FFFF:FFFF (e.g. ARP packet). [PR1073536](#)
- It tries to check allotted power for all the FPCs, here in the CHASSISD_I2CS_READBACK_ERROR logs it shows for the FPCs which are not present in chassis. It just calls i2cs_readback() to read i2c device and fails there as these FPCs? slots are blank and prints those readback errors. Also the errors are harmless: "CHASSISD_I2CS_READBACK_ERROR: Readback error from I2C slave for FPC" Fix: Code to check 'if power has been allotted to this FPC', needs to be executed only if the FPC is present. [PR1075643](#)
- When using the "ping detail" command, the interface number is provided on the output instead of the interface name. [PR1078300](#)
- During a unified in-service software upgrade (ISSU), DHCP control traffic (renew/rebinds) might be dropped on ingress Packet Forwarding Engine. [PR1079812](#)
- When an MX chassis network-services is "enhanced-ip" and an AE is part of a Layer 2 bridge (bridge-domain or VPLS), there is a possibility that an incorrect forwarding path might be installed causing traffic loss. This could happen when first applying the configuration, restarting the system or restarting the line card. [PR1081999](#)
- On MX Series-based platform, the "RPF-loose-mode-discard" feature is not working when configured within a Virtual Router routing instance. The feature is working only when configured in the main instance. [PR1084715](#)
- With MSDPC equipped on BNG, there might be a memory leak in ukernel, which eventually causes MSDPC to crash and restart. [PR1085023](#)
- In Junos OS Releases 13.3R3, 14.1R1, 14.2R1, there is a new feature, an extra TLV term is added to accommodate the default action for the "next-interface" when the corresponding next-interface is down. While doing a unified ISSU from an image without the feature to an image with this feature, all MPCs might crash. [PR1085357](#)
- If there are scaling unicast routes (e.g. 500k) in NG-MVPN VRF, and the provider-tunnel is PIM, when PIM on PE has multiple upstream neighbors and any of them could be its rpf neighbor, performing GRES/NSR Routing Engine switchover might cause multicast traffic loss due to the different view of rpf neighbor between the master Routing Engine and the slave Routing Engine. [PR1087795](#)
- The prompt for SSH password changed in Junos OS Release 13.3, from "user@host's password:" to "Password:". This change breaks the logic in "JUNOS/Access/ssh.pm"

which is located in /usr/local/share/perl/5.18.2/ on Ubuntu Linux, for example.
[PR1088033](#)

- On MX Series router with MPC1/1E, MPC2/2E line cards in a broadband edge environment with scaled (in this case 250K) subscribers, the FPC heap (dynamic memory) utilization increases significantly during an in-service software upgrade (ISSU). [PR1088427](#)
- On MX Series platform with MPC/MIC or T4000 FPC5, TCP session with MS-Interface/AMS-Interface, configuration is not established successfully with the "no-destination-port" or "no-source-port" configuration statements configured under forwarding-options hierarchy level. [PR1088501](#)
- Issue is specific to 64-Bit RPD and config-groups wildcard configuration specific as in the following case: set groups TEST routing-instances <*> routing-options multicast forwarding-cache family inet threshold suppress 200 set routing-instances vrf1 apply-groups TEST set routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 600 With this daemon(rpd) reads suppressed value ?200? (i.e. coming from groups) instead of reading value ?600? from foreground and customer sees unexpected behavior with respect to threshold-suppress. Workaround: They can replace wildcard with actual routing-instance name as in below example: set groups TEST routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 200 set routing-instances vrf1 apply-groups TEST set routing-instances vrf1 routing-options multicast forwarding-cache family inet threshold suppress 600 [PR1089994](#)
- On MX Series router, if ifl (logical interface) is configured with VID of 0 and parent ifd (physical interface) with native-vlan-id of 0, when sending L2 traffic received on the ifl to Routing Engine, the VID 0 will not imposed, causing the frames to get dropped at Routing Engine. [PR1090718](#)
- When an interface on MQ-based FPC is going to link down state, in-flight packet on interface transmit path will be stuck on the interface and never drained until the interface comes up again. As a result, small number of such stacked packets will be sent out when the interface is going to UP state. No other major impact should be seen after those packets are drained. [PR1093569](#)
- On MX2020/2010 router, an SPMB core file will be seen if there are bad XF chips (fabric chip) on SFB, which might trigger Routing Engine/CB switchover. [PR1096455](#)
- In 64-bit Junos OS environment, the Representational State Transfer (REST) API fails to start when configured with "set system services rest ...". [PR1097266](#)
- When a P2MP LSP is added or deleted at ingress LSR, traffic loss is seen to existing sub-LSP(s) at transit LSR which replicates and forwards packet to egress PEs. This issue only affects MX Series based line card. [PR1097806](#)
- The "shared-bandwidth-policer" configuration statement is used to enable configuration of interface-specific policers applied on an aggregated Ethernet bundle to match the effective bandwidth and burst-size to user-configured values. But this feature is broken from Junos OS Release 14.1R1 when "enhanced-ip" is configured on MX Series platform with pure MX Series-based line cards. The bandwidth/burst-size

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

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

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

Routing Policy and Firewall Filters

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

Routing Protocols

- Support for the Pragmatic General Multicast protocol (daemon pgmd) is being phased out from Junos OS. In Junos OS Release 14.2, the CLI is now hidden (although the component is still there and configurable). In Junos OS Release 15.1 the code and its corresponding CLI are removed. [PR936723](#)
- In PIM multicast-only fast reroute (MoFRR) environment, when issuing CLI command "show multicast route extensive" on egress edge router, due to missing null check while showing label information for reverse-path forwarding (RPF) nexthop, an error might be seen in the output of the command. In addition, the routing protocol process (rpd) may crash on the device. [PR983140](#).
- For the pim nbr which is not directly connected (that is, nbr on unnumbered interface, or p2p interface with different subnet), pim join is not able to find the correct upstream nbr which results in join not propagating to the upstream nbr. show command for pim join shows upstream nbr "unknown". Issue is present in the 15.1R1 release. [PR1069896](#)
- In multi-topologies IS-IS scenario, there is huge difference between estimated free bytes and actual free bytes when generating LSP with IPv6 Prefix. It might cause LSP fragment exhaustion. [PR1074891](#).
- There are two issues in the PR: (1) In multicast environment, Incoming interface list (IIF) list has only RPF interface, designated forwarder (DF) winners are not added in the list in backup Routing Engine. (2) "Number of downstream interfaces" in show pim join extensive is not accounting Pseudo-VXLAN interface. [PR1082362](#)
- If a policy statement referred to a routing-table, but the corresponding routing instance is not fully configured (ie. no instance-type), commit such configuration might cause the rpd process to crash. [PR1083257](#).
- With Multicast Source Discovery Protocol (MSDP) and nonstop active routing (NSR) configured on the Protocol Independent Multicast (PIM) sparse-mode rendezvous point (RP), the rpd process might permanently get stuck when multicast traffic received shortly after Routing Engines switchover. [PR1083385](#)

- When there are a number of secondary BGP routes in inet.0, an SNMP walk of inet.0 by the bgp4 MIB can cause a core if the corresponding primary routes are being deleted. [PR1083988](#).
- 1. configure the ospf and ospf3 in all routers 2. configure node protection 3. check for 22.1.1.0 any backup is present 4. enable pplfa all 5. check for 22.1.1.0 any pplfa backup is present through r2 we are not seeing any pplfa backup for 22.1.1.0 [PR1085029](#)
- When BGP route is leaked to a routing-instance and there is an import policy to overwrite the route preference, if damping is also configured in BGP, the BGP routes which were copied to second table cannot be deleted after routes were deleted in master table. This is a day-1 issue. [PR1090760](#)
- When removing BGP Prefix-Independent Convergence (PIC) from the configuration, the expected behavior is that any protected path would become unprotected. But in this case, the multipath entry that contains the protection path (which is supposed to be removed) remains active, until BGP session flaps or the route itself flaps. As a workaround, we can use "commit full" command to correct or to commit. [PR1092049](#)
- In BGP environment, when configuring RIB copy of routes from primary routing table to secondary routing table (for example, by using the CLI command "import-rib [inet.0 XX.inet.0]") and if the second route-table's instance is type "forwarding", due to the BGP routes in secondary routing table may get deleted and not correctly re-created, the routes may be gone on every commit (even commit of unrelated changes). As a workaround, for re-creating the BGP routes in secondary route table, use CLI command "commit full" to make configuration changes. [PR1093317](#)
- In Junos OS Release 9.1 and later, RFC 4893 introduces two new optional transitive BGP attributes, AS4_PATH and AS4_AGGREGATOR. These new attributes are used to propagate 4-byte AS path information across BGP speakers that do not support 4-byte AS numbers. In this case, when AS4_AGGREGATOR attribute (18) is received from a 2-byte AS peer (note AS4_AGGREGATOR attribute is only received when the aggregator has 4-byte AS but this peer only supports 2-byte AS), NSR synchronization with standby Routing Engine would fail, causing session constantly bouncing on standby Routing Engine (hogging CPU). [PR1093615](#)
- The rpd process might crash when resolve-vpn and rib inet.3 are configured under separate levels (BGP global, group and peer). The fix is if anybody configures a family at a lower level, reset the state created by either of configuration statements from higher levels. This behavior conforms with our current behavior of family configuration - which is that any configuration at a lower level is honored and the higher level configuration is reset. [PR1094499](#).
- When BGP routes has multiple protocol nexthops including discard/reject and other IGP nexthops, the discard/reject nexthop will be selected as BGP nexthop, which will cause traffic loss. [PR1096363](#)
- When polling SNMP OID isisPacketCounterTable 1.3.6.1.2.1.138.1.5.3, the rpd process might crash. [PR1101080](#).
- When the IS-IS configurations have been removed, the IS-IS LSDB contents get flushed. If at the same time of this deletion process, there is an SPF execution (that is, try to

access the data structures at same time when/a fraction of seconds after freeing its content), routing protocol process (rpd) crash occurs. [PR1103631](#)

Services Applications

- When an MX Series router configured as an LNS sends an Access-Request message to RADIUS for an LNS subscriber, the LNS now includes the Called-Station-ID-Attribute when it receives AVP 21 in the ICRQ message from the LAC. [PR790035](#)
- When an MS-DPC PIC reboots due to a crash or manual intervention, it might get stuck in a booting loop if the MS-DPC up-time is more than 49 days and 17 hours. After 5 consecutive boot failures, the MS-DPC PIC will go offline automatically and give the following error message: [15:21:22.344 LOG: Err] ICHIP(0): SPI4 Training failed while waiting for PLL to get locked, ichip_sr a_spi4_r x_snk_init_s tatus_clk [15:21:22.344 LOG: Err] CMSP C:I-Chip(0) SPI4RxSinkinitstatusclockfailed, cmsdpc_spi4_init [15:21:22.344 LOG: Err] CMX: I(0) ASIC SPI4 init failed [15:21:22.379 LOG: Err] Node for service control ifl 68, is already present [15:21:23.207 LOG: Err] ASER0 SPI-4 XLR source core OOF did not go low in 20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 stop src train failed! [15:21:23.208 LOG: Err] ASER0 XLR SPI-4 sinkcoreDPAINcompletein20ms. [15:21:23.208 LOG: Err] ASER/XLR0 spi4 sink core init failed! [15:21:24.465 LOG: Err] ICHIP(0): SPI4 Stats Unexpected 2'b 11 Error, isra_spi4_p arse_panic_err ors [15:21:24.465 LOG: Err] ICHIP(0): SPI4 Tx Lost Sync Error, isra_spi4_p arse_panic_err ors . In order to recover from this state, the whole MS-DPC needs to be rebooted. [PR828649](#)
- In IPsec environment, after performing the Routing Engine switchover (for example, performing Graceful Routing Engine Switchover) or chassis reboot (that is, whole device is powered down and powered UP again), due to the key management daemon (kmd) may be launched before the Routing Engine mastership is finalized, it may stop running on the new master Routing Engine. [PR863413](#)
- On MX Series platform, when using the MS-DPC with MPSDK to support Captive Portal Content Delivery (cpcd) service, the MAC might get stuck on the FPC due to processing the high rate of packets (for example, 5kpps HTTP traffic). In addition, reloading the affected FPC might only temporarily resolve the issue while it will appear again once scaling up. [PR1037143](#)
- In CG-NAT or statefull firewall environment, due to a null pointer check bug, the MS-DPC might crash every few hours. Note that this is a regression issue. [PR1079981](#)
- The crash happens if in a http flow, the flow structure is allocated at a particular memory region. There is no workaround but the chances of hitting this issue are very low [PR1080749](#)
- On Layer 2 Tunnel Protocol (L2TP) network server (LNS), during L2TP session establishment, when receiving Incoming-Call-Connected (ICCN) messages with Last Sent LCP CONFREQ Attribute Value Pair (AVP) but without Initial Received LCP CONFREQ and Last Received LCP CONFREQ AVPs, the jl2tpd process might crash. [PR1082673](#)
- On Trivial File Transfer Protocol (TFTP) Application Layer Gateway (ALG) with NAT translation type "dynamic-nat44" configured, MS-DPC/MS-MPC/MS-MIC might crash when processes the TFTP packets. [PR1091179](#)

- On M Series platform, in Layer 2 Tunneling Protocol (L2TP) network server (LNS) environment, not all attributes (Missing NAS-Identifier, NAS-Port-Type, Service-Type, Framed-Protocol attributes) within Accounting-Request packet are sending to the RADIUS server. [PR1095315](#)
- If MS-DPC is used in CG-NAT environment, in a very rare condition, when the MS-DPC tries to delete a NAT mapping entry (e.g. entry timeout), error might occur and the MS-DPC might get rebooted and then generate a core file. [PR1095396](#)
- Some values of MIB object jnxSrcNatStatsEntry might be doubled when AMS (or rsp) interface and NAT are configured together. [PR1095713](#)

Software Installation and Upgrade

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

Subscriber Access Management

- In subscriber management environment, after deactivating a service with Change of Authorization (CoA) dynamic requests, if the Acct-Stop response is not received, the Broadband Network Gateway (BNG) will send CoA NAK message when the same service is activated again. The authd process crash will be observed and some sessions are stuck and cannot be terminated after terminating sessions. [PR1004478](#)
- The authd process memory leaks slowly when subscribers login and logout, which eventually leads the process to crash and generate a core file. [PR1035642](#)
- On MX Series routers, the generic authentication service process (authd) may fail to send Acct-off message to the RADIUS server. This is because management daemon (mgd) might not notify the authd prior to executing system reboot or system shutdown. Also, the authd might fail to generate the Acct-off message as well when it is terminated and there are no active subscribers. [PR1053044](#)
- In subscriber management environment with Remote Authentication Dial In User Service (RADIUS) server configured, when performing scaling subscribers login/logout, the device may stuck in RADIUS communication. [PR1070468](#)
- In subscriber management environment, when dual-stack service is activated by the Change of Authorization (CoA) request from the Radius Server, both families will be activated in the same profile response. Due to a software defect, the service accounting session id is not generated properly and the Service Accounting Messages and Interim-updates failed to be sent out. [PR1071093](#)
- Subscriber is not coming up when CISCO AVPair VSA value is returned in Radius ACCESS-ACCEPT packets in certain scenarios. [PR1074992](#)
- A CoA Request containing LI attributes cannot contain any non-LI service activations, de-activations or variable modifications. [PR1079036](#)

- If authentication-order is configured as none under access profile and domain-name servers (DNS) are configured locally under access profile, then the subscriber will login but will not get DNS addresses which were configured locally. [PR1079691](#)
- In scaled DHCP subscribers environment, the authd process might crash and generate a core file after clearing DHCP binding or logout subscribers. [PR1094674](#)

User Interface and Configuration

- Due to a change in an existing PR, group names in the configuration must be a string of alphanumericals, dashes or underscores. There is not workaround other than following the group name instructions. [PR1087051](#)

VPNs

- Problem, trigger and symptom: On dual Routing Engines, if mvpn protocol itself is not configured, and non stop routing is enabled, the show command "show task replication" on master Routing Engine will list MVPN protocol even though it is not configured. Other than the misleading show output which may be slightly confusing to the user/customer, there is no functional impact due to this issue as such. There is no workaround available. [PR1078305](#)
- In PIM Draft-Rosen Multicast VPN (MVPN) environment, in a setup where active C-RP, standby C-RP, C-receivers, C-source are located in different VPN site of MVPN instance, once the link to active C-RP is flapped, PE which connects to C-receivers would send (*,g) join and (s,g,rpt) prune towards standby C-RP, when the PE which connects to standby C-RP receives the (*,g) join and (s,g, rpt) prune over mt-, it ends up updating the (s,g) forwarding entry with mt- as downstream, which is already the incoming interface (IIF). This creates a forwarding loop due to missing check if IIF is same as OIF when PIM make-before-break (MBB) join load-balancing feature is enabled and as a result traffic gets looped back into the network. Loop once formed will remain at least for 210 seconds till the delayed prune timer expires. After this, IIF is updated to the interface towards standby C-RP finally. [PR1085777](#)
- In NG-MVPN spt-only mode with a PE router acts as the rendezvous point (RP), if there are only local receivers, the unnecessary multicast traffic continuously goes to this RP and dropped though it is not in the shortest-path tree (SPT) path from source to receiver. [PR1087948](#)
- When there are more than 2000 outgoing interfaces (OIFs) for a same multicast group on MVPN egress PE, the multicast forwarding entries installed by MVPN might have duplicated OIFs and resulting in duplicated traffic. [PR1095877](#)
- In Internet multicast over an MPLS network by using next-generation Layer 3 VPN multicast (NG-MVPN) environment, when rib-groups are configured to use inet.2 as RPF rib for Global Table Multicast (GTM, internet multicast) instance, the ingress PE may fail to add P-tunnel as downstream even after receiving BGP type-7 routes. In addition, this issue only affects GTM. [PR1104676](#)

- See Also**
- [New and Changed Features on page 80](#)
 - [Changes in Behavior and Syntax on page 131](#)

- [Known Behavior on page 165](#)
- [Known Issues on page 170](#)
- [Documentation Updates on page 312](#)
- [Migration, Upgrade, and Downgrade Instructions on page 320](#)
- [Product Compatibility on page 329](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R6 documentation for the M Series, MX Series, and T Series.

- [Adaptive Services Interfaces Feature Guide for Routing Devices on page 312](#)
- [Advanced Subscriber Management Provisioning Guide on page 313](#)
- [Broadband Subscriber Sessions Feature Guide on page 313](#)
- [Broadband Subscriber VLANs and Interfaces Feature Guide on page 314](#)
- [High Availability Feature Guide on page 314](#)
- [IPv6 Neighbor Discovery Feature Guide for Routing Devices on page 315](#)
- [Monitoring, Sampling, and Collection Services Interfaces Feature Guide for Routing Devices on page 315](#)
- [MPLS Applications Feature Guide for Routing Devices on page 316](#)
- [Overview for Routing Devices on page 317](#)
- [Release Notes on page 317](#)
- [Routing Policies, Firewall Filters, and Traffic Policers Feature Guide for Routing Devices on page 318](#)
- [Security Services Administration Guide for Routing Devices on page 318](#)
- [Standards Reference on page 318](#)
- [Subscriber Management Access Network Guide on page 318](#)
- [Subscriber Management Provisioning Guide on page 319](#)
- [Tunnel and Encryption Services Interfaces on page 319](#)
- [User Access and Authentication Guide for Routing Devices on page 319](#)
- [VPNs Library for Routing Devices on page 319](#)

Adaptive Services Interfaces Feature Guide for Routing Devices

- In the topic “Inline 6rd and 6 to 4 Configuration Guidelines”, the next-to-last bullet should state:

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

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



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

- The descriptions in the “Options” section of the IPsec **protocol** statement at the `[edit services ipsec-vpn ipsec proposal proposal-name]` and `[edit services ipsec-vpn rule rule-name term term-name then manual direction direction]` hierarchy levels fail to state that the **ah** and **bundle** options are not supported on MS-MPCs and MS-MICs on MX Series routers.

Advanced Subscriber Management Provisioning Guide

- The “Example: Configuring HTTP Redirect Services on the Routing Engine” topic shows an incorrectly formatted redirect URL, `http://www.google.com?=%dest-url%`. The correct format is `http://www.example.com/url=%dest-url%`.

Broadband Subscriber Sessions Feature Guide

- The “enhanced-policer” topic erroneously states that when you commit a configuration that includes this statement, the CLI displays a warning that FPCs must be restarted for it to take effect, and prompts you to proceed with a restart. No such warning or prompt is displayed; instead, a warning message is logged that states that the enhanced policer is enabled on FPCs only after they are restarted.
- The following topics erroneously include information about the Ignore-DF-Bit VSA (26-70): “RADIUS Attributes and Juniper Networks VSAs Supported by the AAA Service Framework”, “Juniper Networks VSAs Supported by the AAA Service Framework”, and “AAA Access Messages and Supported RADIUS Attributes and Juniper Networks VSAs for Junos OS”. Junos OS does not support VSA 26-70.

Some versions of the RADIUS dictionary file also erroneously list 26-70 as supported by the Junos OS.

- The following topics indicate that you can configure an MX Series router to maintain a DHCP subscriber in the event the subscriber interface is deleted:
 - “Subscriber Binding Retention During Interface Delete Events”
 - “Configuring the Router to Maintain DHCP Subscribers During Interface Delete Events”
 - “Verifying and Managing the DHCP Maintain Subscribers Feature”
 - “interface-delete (Subscriber Management or DHCP Client Management)”
 - “maintain-subscriber”
 - “subscriber-management (Subscriber Management)”

This feature is not supported on MX Series routers running Junos OS Release 15.1R4 or later with enhanced subscriber management enabled.

- The Broadband Subscriber Sessions Feature Guide did not report the single session DHCP dual-stack feature, which enables the use of only a single session for authentication rather than the three sessions required for the traditional dual-stack configuration. See the description of this feature in [“New and Changed Features” on page 80](#).

Broadband Subscriber VLANs and Interfaces Feature Guide

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

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

High Availability Feature Guide

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

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

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

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

- The “Nonstop Active Routing System Requirements” topic should include the **inet-mvpn** and **inet6-mvpn** protocol families for BGP in the list of supported family types. The topic previously documented that NSR supports next-generation MVPN starting with Junos OS 14.1R1, but didn't include the specific names of the next-generation MVPN protocol families in the list.
- The topic “Improving the Convergence Time for VRRP” failed to include the following information:

- Disable duplication address detection for IPv6 interfaces—Duplicate address detection is a feature of the Neighbor Discovery Protocol for IPv6. Duplicate address detection is enabled by default and determines whether an address is already in use by another node. When duplicate address detection is enabled, convergence time is high after an IPv6 interface that has been configured for VRRP tracking comes up. To disable duplicate address detection, include the **ipv6-duplicate-addr-detection** statement at the **[edit system internet-options]** hierarchy level. To disable duplicate address detection only for a specific interface, include the **dad-disable** statement at the **[edit interfaces interface-name unit logical-unit-number family inet6]** hierarchy level.

IPv6 Neighbor Discovery Feature Guide for Routing Devices

- The *Secure Neighbor Discovery Guide for Routing Devices* is merged with the *IPv6 Neighbor Discovery Feature Guide for Routing Devices*. We have consolidated these guides and restructured the content in a linear format. The new seamless guide provides related information in a single location for easy navigation and faster access.

[See [IPv6 Neighbor Discovery Feature Guide for Routing Devices](#).]

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

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

- The Options section for the **flow-export-rate** statement under the hierarchy **[edit forwarding-options sampling instance instance-name family inet output inline-jflow]** did not include the default value. The default value is:

Default: 1 for each Packet Forwarding Engine on the FPC to which the sampling instance is applied.

- The following topics fail to state that for passive monitoring on MX Series routers with MPCs, the **pop-all-labels** statement at the **[edit interfaces interface-name]** hierarchy level pops all labels by default, and the **required-depth** statement is ignored.
 - “pop-all-labels”
 - “required-depth”
 - “Enabling Passive Flow Monitoring”
- The “Configuring RPM Timestamping” topic failed to mention that RPM timestamping is also supported on the MS-MPCs and MS-MICs on MX Series routers.
- The description for the **max-packets-per-second**, **maximum-packet-length**, and **run-length** statements at the **[edit forwarding-options sampling instance instance-name input]** hierarchy level failed to include the following:



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

- The default value for the `ipv6-flow-table-size` statement at the `[edit chassis fpc slot-number inline-services ipv6 flow-table-size]` hierarchy level should state the following:
 "If the number of units is not specified, 1024 flow entries are allocated for IPv6."
- The topics "Real-Time Performance Monitoring Services Overview" and "Configuring RPM Probes" failed to state that RPM is not supported on logical systems.
- The following topics should state that the `test-interval` statement at the `[edit services rpm probe owner test test-name]` hierarchy level has a range from 0 through 86400 seconds, and that a value of 0 seconds causes the RPM test to stop after one iteration:
 - "Configuring RPM Probes"
 - "test-interval"
 - "Configuring BGP Neighbor Discovery Through RPM"

MPLS Applications Feature Guide for Routing Devices

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

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

Overview for Routing Devices

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

Release Notes

- The release notes for the following Junos OS releases incorrectly included a new feature that reported support for VLAN demux interfaces on MS-DPCs:
 - 15.1R1 Release Notes
 - 15.1R2 Release Notes
 - 15.1R3 Release Notes
 - 15.1R4 Release Notes
 - 15.1R5 Release Notes

VLAN demux interfaces are not supported on MS-DPCs in Junos OS Release 15.1R1 and later releases. VLAN demux interfaces in those releases require enhanced subscriber management. Enhanced subscriber management does not support MS-DPCs.

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

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

Security Services Administration Guide for Routing Devices

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

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

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

Standards Reference

- The *Supported Network Management Standards* topic incorrectly states that Junos OS supports `mplsL3VpnIfConfTable` as part of compliance with RFC 4382, MPLS/BGP Layer 3 Virtual Private Network (VPN) MIB. Junos OS does not support this table.

Subscriber Management Access Network Guide

- The “Configuring a Pseudowire Subscriber Logical Interface Device” and “anchor-point (Pseudowire Subscriber Interfaces)” topics have been updated to state that you cannot dynamically change an anchor point that has active pseudowire devices stacked above it. Both topics describe the steps to follow when you must change such an anchor point.
- The following topics have been updated to reflect a change in recommendation for use of the **access-internal** statement: “Access and Access-Internal Routes for Subscriber Management”, “Configuring Dynamic Access Routes for Subscriber Management”,

“Access (Dynamic Access Routes)”, and “Access-internal (Dynamic Access-Internal Routes)”.

Starting in Junos OS Release 15.1, we recommend that you use only access routes for framed route support. We recommend that you do not use access-internal routes. If the RADIUS Framed-Route attribute (22) or Framed-IPv6-Route attribute [99] does not specify the next-hop gateway—as is common—the variable representing the next-hop, `$junos-framed-route-nexthop`, is automatically resolved. If you configure the **access-internal** statement in the dynamic profile, it is ignored.

Subscriber Management Provisioning Guide

- The topic “Configuring Address-Assignment Pool Linking” states that when you link multiple address-assignment pools, a secondary pool is used only when the primary address-assignment pool is fully allocated. However, once the router switches to a pool other than the primary, it continues using that pool even when addresses are available again in the primary pool.

Tunnel and Encryption Services Interfaces

- The topic “Configuring Tunnel Interfaces on MX Series Routers” incorrectly states that bandwidth rates of 20 gigabits per seconds and 40 gigabits per second require use of a 100-Gigabit Ethernet Modular Port Concentrator and 100-Gigabit CFP MIC. The MPC4E, MPC5E, and MPC6E also support 20 and 40 gigabits per second.

User Access and Authentication Guide for Routing Devices

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

VPNs Library for Routing Devices

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

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

- See Also**
- [New and Changed Features on page 80](#)
 - [Changes in Behavior and Syntax on page 131](#)

- [Known Behavior on page 165](#)
- [Known Issues on page 170](#)
- [Resolved Issues on page 192](#)
- [Migration, Upgrade, and Downgrade Instructions on page 320](#)
- [Product Compatibility on page 329](#)

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the M Series, MX Series, and T Series. Upgrading or downgrading Junos OS can take several minutes, depending on the size and configuration of the network.

Starting with Junos OS Release 15.1, in some of the devices, FreeBSD 10.x is the underlying OS for Junos OS instead of FreeBSD 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.1 remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see [Understanding Junos OS with Upgraded FreeBSD](#).



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

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

Platform	FreeBSD 6.1-based Junos OS	FreeBSD 10.x-based Junos OS
M7i, M10i, M120, M320	YES	NO
MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES
T640, T1600, T4000, TX Matrix, TX Matrix Plus	YES	NO

- [Basic Procedure for Upgrading to Release 15.1 on page 321](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 10.x\) on page 322](#)
- [Upgrading from Junos OS \(FreeBSD 6.1\) to Junos OS \(FreeBSD 6.1\) on page 323](#)
- [Upgrade and Downgrade Support Policy for Junos OS Releases on page 325](#)
- [Upgrading a Router with Redundant Routing Engines on page 325](#)

- [Upgrading Juniper Network Routers Running Draft-Rosen Multicast VPN to Junos OS Release 10.1 on page 326](#)
- [Upgrading the Software for a Routing Matrix on page 327](#)
- [Upgrading Using Unified ISSU on page 328](#)
- [Downgrading from Release 15.1 on page 329](#)

Basic Procedure for Upgrading to Release 15.1

When upgrading or downgrading Junos OS, always use the **jinstall** package. Use other packages (such as the **jbundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#) and [Upgrading Junos OS with Upgraded FreeBSD](#).



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

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

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

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

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



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

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

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



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

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-15.1R6.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-15.1R6.9-signed.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.



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

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

Products impacted: All M Series routers, all T Series routers, MX80, and MX104.



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

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

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



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

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

```
user@host> request system software add validate reboot  
source/jinstall-15.1R6.9-domestic-signed.tgz
```

- All other customers, use the following command:

```
user@host> request system software add validate reboot  
source/jinstall-15.1R6.9-export-signed.tgz
```

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

- **/pathname**—For a software package that is installed from a local directory on the router.

- For software packages that are downloaded and installed from a remote location:
 - `ftp://hostname/pathname`
 - `http://hostname/pathname`
 - `scp://hostname/pathname` (available only for Canada and U.S. version)

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 15.1 **jinstall** package, you cannot issue the `request system software rollback` command to return to the previously installed software. Instead, you must issue the `request system software add validate` command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or from Junos OS Release 15.1 to Release 16.2. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

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

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

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

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



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

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

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



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

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

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

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

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

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



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

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

Upgrading the Software for a Routing Matrix

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

- A minimum of free disk space and DRAM on each Routing Engine. The software upgrade will fail on any Routing Engine without the required amount of free disk space and DRAM. To determine the amount of disk space currently available on all Routing Engines of the routing matrix, use the CLI **show system storage** command. To determine the amount of DRAM currently available on all the Routing Engines in the routing matrix, use the CLI **show chassis routing-engine** command.
- The master Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC are all re0 or are all re1.
- The backup Routing Engines of the TX Matrix or TX Matrix Plus router (SCC or SFC) and all LCCs connected to the SCC or SFC are all re1 or are all re0.

- All master Routing Engines in all routers run the same version of software. This is necessary for the routing matrix to operate.
- All master and backup Routing Engines run the same version of software before beginning the upgrade procedure. Different versions of Junos OS can have incompatible message formats especially if you turn on GRES. Because the steps in the process include changing mastership, running the same version of software is recommended.
- For a routing matrix with a TX Matrix router, the same Routing Engine model is used within a TX Matrix router (SCC) and within a T640 router (LCC) of a routing matrix. For example, a routing matrix with an SCC using two RE-A-2000s and an LCC using two RE-1600s is supported. However, an SCC or an LCC with two different Routing Engine models is not supported. We suggest that all Routing Engines be the same model throughout all routers in the routing matrix. To determine the Routing Engine type, use the CLI **show chassis hardware | match routing** command.
- For a routing matrix with a TX Matrix Plus router, the SFC contains two model RE-DUO-C2600-16G Routing Engines, and each LCC contains two model RE-DUO-C1800-8G or RE-DUO-C1800-16G Routing Engines.



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

To upgrade the software for a routing matrix:

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

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

Upgrading Using Unified ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified ISSU, see the [High Availability Feature Guide for Routing Devices](#).

For information on ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\) Web application](#).

Downgrading from Release 15.1

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



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

For more information, see the [Installation and Upgrade Guide](#).

- See Also**
- [New and Changed Features on page 80](#)
 - [Changes in Behavior and Syntax on page 131](#)
 - [Known Behavior on page 165](#)
 - [Known Issues on page 170](#)
 - [Resolved Issues on page 192](#)
 - [Documentation Updates on page 312](#)
 - [Product Compatibility on page 329](#)

Product Compatibility

- [Hardware Compatibility on page 329](#)

Hardware Compatibility

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

To determine the features supported on M Series, MX Series, and T Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:

<https://pathfinder.juniper.net/feature-explorer/>

- See Also**
- [New and Changed Features on page 80](#)
 - [Changes in Behavior and Syntax on page 131](#)
 - [Known Behavior on page 165](#)

- [Known Issues on page 170](#)
- [Resolved Issues on page 192](#)
- [Documentation Updates on page 312](#)
- [Migration, Upgrade, and Downgrade Instructions on page 320](#)

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 15.1R6 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

- [New and Changed Features on page 331](#)
- [Changes in Behavior and Syntax on page 340](#)
- [Known Behavior on page 343](#)
- [Known Issues on page 343](#)
- [Resolved Issues on page 346](#)
- [Documentation Updates on page 360](#)
- [Migration, Upgrade, and Downgrade Instructions on page 361](#)
- [Product Compatibility on page 365](#)

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 15.1R6 for the PTX Series.

- [High Availability and Resiliency \(HA\) on page 332](#)
- [Interfaces and Chassis on page 332](#)
- [IPv6 on page 333](#)
- [Junos OS XML API and Scripting on page 333](#)
- [Management on page 334](#)
- [MPLS on page 335](#)
- [Routing Protocols on page 335](#)
- [Software Licensing on page 336](#)
- [User Interface and Configuration on page 339](#)
- [VPNs on page 339](#)

High Availability and Resiliency (HA)

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

Interfaces and Chassis

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

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

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

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

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

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

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

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

IPv6

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

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

Junos OS XML API and Scripting

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

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

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

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

Management

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

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

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

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

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

MPLS

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

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

Routing Protocols

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

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

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

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

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

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

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

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

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

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

Software Licensing

- **Licensing enhancements (PTX Series)**—Starting with Junos OS Release 15.1R1, licensing enhancements on PTX Series routers enable you to configure and delete license keys in a Junos OS CLI configuration file. The license keys are validated and installed after a successful commit of the configuration file. If a license key is invalid, the commit fails and issues an error message. You can configure individual license keys or multiple license keys by issuing Junos OS CLI commands or by loading the license key configuration contained in a file. All installed license keys are stored in the `/config/license/` directory.

To install an individual license key in the Junos OS CLI, issue the **set system license keys key *name*** command, and then issue the **commit** command.

For example:

```
[edit]
root@switch# set system license keys key "JUNOS_TEST_LIC_FEAT aeaqeb qbmqds
qwwsxe okyvou 6v57u5 zt6ie6 uwl3zh assvnu e2ptl5 soxawy vtfh7k axwnno m5w54j
6z"
root@switch# commit
commit complete
```

To verify that the license key was installed, issue the **show system license** command.

For example:

```
root@switch> show system license
License usage:

```

	Licenses	Licenses	Licenses	Expiry
Feature name	used	installed	needed	
sdk-test-feat1	0	1	0	
permanent				

```

Licenses installed:
License identifier: JUNOS_TEST_LIC_FEAT
License version: 2
Features:
  sdk-test-feat1 - JUNOS SDK Test Feature 1
  permanent
```

To install multiple license keys in the Junos OS CLI, issue the **set system license keys key *name*** command, and then issue the **commit** command.

For example:


```
[edit]
root@switch# set system license keys key "key_1"
set system license keys key "key_2"
set system license keys key "key_2"
set system license keys key "key_4"
root@switch# commit
commit complete
```

To verify that the license key was installed, issue the **show system license** command.

To install an individual license key configuration in a file, issue the **cat** command:

For example:

```
[edit]
root@switch%cat license.conf
system {
  license {
    keys {
      key "JUNOS_TEST_LIC_FEAT aeaqeb qbmqds qwwsxe okyvou 6v57u5 zt6ie6
uw13zh assvnu e2pt15 soxawy vtfh7k axwnno m5w54j 6z";
    }
  }
}
```

Load and merge the license configuration file.

For example:

```
[edit]
root@switch# load merge license.conf
load complete
```

Issue the **show | compare** command to see the configuration, and then issue the **commit** command.

For example:

```
[edit]
root@switch# show | compare
[edit system]
+ license {
+   keys {
+     key "JUNOS_TEST_LIC_FEAT aeaqeb qbmqds qwwsxe okyvou 6v57u5
zt6ie6 uw13zh assvnu e2pt15 soxawy vtfh7k axwnno m5w54j 6z";
+   }
+ }
[edit]
root@switch# commit
```

To verify that the license key was installed, issue the **show system license** command.

For example:

```
root@switch> show system license
License usage:
Lenses      Lenses      Lenses      Expiry
```

Feature name	used	installed	needed
sdk-test-feat1	0	1	0
permanent			

Licenses installed:
 License identifier: JUNOS_TEST_LIC_FEAT
 License version: 2
 Features:
 sdk-test-feat1 - JUNOS SDK Test Feature 1
 permanent

To install multiple license keys in a file, issue the **cat** command:

For example:

```
[edit]
root@switch%cat license.conf
system
{
  license
  {
    keys
    {
      key "key_1"
      key "key_2"
      key "key_3"
      ...
      key "key_n"
    }
  }
}
```

Load and merge the license configuration file, and then issue the **commit** command.

For example:

```
[edit]
root@switch# load merge license.conf
load complete
[edit]
root@switch# commit
```

To verify that the license key was installed, issue the **show system license** command.

You can also delete or deactivate individual and multiple license keys in the Junos OS CLI by issuing the **delete system license keys** or **deactivate system license keys** commands. Do not use the **request system license delete** command to delete the license keys.

For example, to issue the **delete system license keys** command:

```
[edit]
root@switch# delete system license keys
root@switch# commit
```

User Interface and Configuration

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

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

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

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



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

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

VPNs

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

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

See Also • [Changes in Behavior and Syntax on page 340](#)

- [Known Behavior on page 343](#)
- [Known Issues on page 343](#)
- [Resolved Issues on page 346](#)
- [Documentation Updates on page 360](#)
- [Migration, Upgrade, and Downgrade Instructions on page 361](#)
- [Product Compatibility on page 365](#)

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 15.1R6 for the PTX Series.

- [High Availability \(HA\) and Resiliency on page 340](#)
- [IPv6 on page 341](#)
- [Junos OS XML API and Scripting on page 341](#)
- [Management on page 341](#)
- [Network Management and Monitoring on page 341](#)
- [Routing Policy and Firewall Filters on page 342](#)
- [Routing Protocols on page 342](#)
- [User Interface and Configuration on page 342](#)

High Availability (HA) and Resiliency

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

[See [request chassis routing-engine master](#).]
- **GRES readiness is part of `show system switchover output` (PTX Series)**—Starting in Junos OS Release 15.1, switchover readiness status is reported as part of the output for operational mode command **show system switchover**.

IPv6

- **IPv6 addresses with padded zeros in MIC or MS-MPC system log messages (PTX Series)**—Starting with Junos OS Release 15.1R2, all system log messages originating from MIC or MS-MPC line cards display padded zeros in IPv6 addresses to make them compatible with MS-DPC line cards. Earlier, the system log messages from MIC or MS-MPC line cards displayed IPv6 addresses with ":::" instead of padded zeros.

Junos OS XML API and Scripting

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

Management

- **Support for status deprecated statement in YANG modules (PTX Series)**—Starting with Junos OS Release 15.1R5, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

Network Management and Monitoring

- **Enhancement for SONET interval counter (PTX Series)**—Starting with Junos OS Release 15.1R3, only the **Current Day Interval Total** output field in the **show interfaces interval** command for SONET interfaces is reset after 24 hours. In addition, the **Previous Day Interval Total** output field displays the last updated time in **hh:mm**.
[See [show interfaces interval](#).]
- **New 64-bit counter of octets for interfaces (PTX Series)**—Starting with Release 15.1R3, Junos OS supports two new Juniper Networks enterprise-specific Interface MIB Extension objects—**ifHCIn1SecOctets** and **ifHCOut1SecOctets**—that act as 64-bit counters of octets passing through an interface.

Routing Policy and Firewall Filters

- **Support for logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol (PTX Series)**—Starting with Junos OS Release 15.1R4, you can configure logical queue-depth in the Packet Forwarding Engine for IP options packets for a given protocol. The queue-depth indicates the number of IP options packets that can be enqueued in the Packet Forwarding Engine logical queue, beyond which it would start dropping the packets.

Routing Protocols

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

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

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

User Interface and Configuration

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

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

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

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

- **New warning message for the configuration changes to extend-size (PTX Series)**—Starting with Junos OS Release 15.1R2, any operation on the **system configuration-database extend-size** configuration statement such as **deactivate**, **delete**, or **set**, generates the following warning message:

Change in 'system configuration-database extend-size' will be effective at next reboot only.

See Also • [New and Changed Features on page 331](#)

- [Known Behavior on page 343](#)
- [Known Issues on page 343](#)
- [Resolved Issues on page 346](#)
- [Documentation Updates on page 360](#)
- [Migration, Upgrade, and Downgrade Instructions on page 361](#)
- [Product Compatibility on page 365](#)

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R6 for the PTX Series.

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

- [System Logging on page 343](#)

System Logging

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

- See Also**
- [New and Changed Features on page 331](#)
 - [Changes in Behavior and Syntax on page 340](#)
 - [Known Issues on page 343](#)
 - [Documentation Updates on page 360](#)
 - [Resolved Issues on page 346](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 361](#)
 - [Product Compatibility on page 365](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R6 for the PTX Series.

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

- [General Routing on page 344](#)
- [Infrastructure on page 345](#)
- [Interfaces and Chassis on page 345](#)
- [MPLS on page 345](#)

- [Platform and Infrastructure on page 345](#)
- [Routing Protocols on page 345](#)

General Routing

- It is reported that on PTX Series platforms, when the firewall filter is configured on the loopback interface of the device, all the FPCs on the device might continuously crash and be unstable because of bad error handling or a NULL pointer. Because the issue is not reproducible, the trigger of the issue is unclear. [PR996749](#)
- On PTX Series platforms, if there are scaling configurations (for example, 5,000 routes and each of them has 64 ECMP paths configured) on a single interface and an L2 rewrite profile is applied for the interface, the FPC might crash when deactivating and then activating the CoS configuration of the interface. [PR1096958](#)
- Starting in Junos OS Release 15.1, there is 1.3 seconds worth of packet loss for I2ck traffic entering or exiting the router during unified ISSU switchover on a PTX5000 with broadway chipset FPCs. Prior to Junos OS Release 15.1, the packet loss was around 0.5 seconds. [PR1102649](#)
- OK LED on the Control Board (CB) does not light up in the following scenarios. 1. When only one RE/CB is inserted on the router and it is powered on. The Master LED is glowing up but not the OK LED. 2. The problem occurs when the test is done with RE/CB on either slot 0 or slot 1, leaving the other slot empty. 3. When both slots have the RE/CB inserted, the problem is not seen. The master and OK LED light up on the master CB, and the OK LED on the backup CB lights up. Engineering is debugging this cosmetic issue. [PR1115148](#)
- PTX 100 GbE-LR4 interfaces might flap when the reference clock switches over from "line clock" to "holdover" initiated by offlining the PIC on which the "line clock" sources reside. When PTX Series router uses "line clock" sources and when it does not have any external clocks from BITS-a or BITS-b, it brings down the "line clock" and the reference clock is switched from "line clock" to "holdover". This reference clock transition might cause a large clock phase-shift in the 100GbE-LR4 CFP modules, and this phase-shift might cause the output optical pulse waveform distortion on the 100 GbE-LR4 interfaces. Hence, it results in interface flap. This issue cannot be fixed by software because of hardware limitation. [PR1130403](#)
- On PTX platforms, when the bits external clock is down or up, the incorrect SNMP trap, jnxFruRemoval(CB), is generated. jnxExtSrcLockAcquired is the correct one. correct trap: Name: "jnxExtSrcLockAcquired" OID: "1.3.6.1.4.1.2636.4.2.5" Incorrect trap: Name: "jnxFruRemoval" OID: "1.3.6.1.4.1.2636.4.1.5" [PR1195686](#)
- Starting in Junos OS Release 15.1, after setting the **set system syslog user** command and rebooting the router, some error log messages appear on the screen. The error messages look like a cosmetic issue only. [PR1223518](#)

Infrastructure

- Multiple negative tests such as a restarting routing or chassis-control might cause the router to reboot. [PR1077428](#)

Interfaces and Chassis

- On dual Routing Engine platforms, when adding the logical interfaces (IFLs) and committing, the device control process (dcd) on the backup Routing Engine might fail to process the configuration and keep it in the memory. In some cases, the memory of the dcd might keep increasing on the backup Routing Engine. [PR1014098](#)
- Configuring ODU FRR under otn-options for the 2x100G DWDM PIC is an unsupported command on the PTX Series router attempting such a configuration could result in an FPC crash and restart. [PR1038551](#)

MPLS

- On a point-to-multipoint (P2MP) MPLS LSP transit router with nonstop active routing (NSR) enabled, when the RSVP refresh reduction feature is enabled and LSP link protection is configured on all interfaces, slight P2MP traffic loss might be seen after the graceful Routing Engine switchover (GRES). [PR1023393](#)
- In an RSVP-based P2MP scenario, if a sub-LSP switchover is used to bypass LSP caused by a PIC offline, when a new sub-LSP is established using setup-protection, deletion of the old sub-LSP might result in deletion of both sub-LSPs. [PR1132050](#)

Platform and Infrastructure

- Parity errors might be seen in preclassifier engines. Such errors are not reported. As such, these packets will be silently discarded. Common symptoms are an increase in value of the "Input DA rejects" counter on the affected interface, or input packets are misclassified. [PR1059137](#)

Routing Protocols

- When you have two paths for the same route, the route gets pointed to Unicast NH which in turn gets pointed to two separate Unicast NHs. The route is determined by OSPF and BFD enabled on one of the paths, which runs through an l2circuit path. When the link on the l2circuit is cut, the link flap is informed by BFD as well as through OSPF LSAs. Ideally the BFD should inform the link-down event before the OSPF LSA. Instead, the OSPF LSAs update the current event a second before BFD. As a result, the route does point to a new Unicast NH with the weights swapped. However, the Unicast NH for which the L3 link is down gets added to the Unicast NH, the BFD assumes the link to be up and updates the weights inappropriately, resulting in traffic loss. Once the BFD link-down event is processed at OSPF protocol level, now the route points to only Unicast NH traffic flows through the currently active link. During FRR, the traffic outage is less than a second. This outage can be avoided if the BFD keepalive intervals are maintained around 50 ms with multiplier of 3 as opposed to 100 ms with a multiplier of 3. [PR1119253](#)

- In a multicast environment, when the rendezvous point (RP) is a first-hop router (FHR) with MSDP peers, when the rpf interface on the RP is changed to an MSDP-facing interface, traffic loss is seen. The loss occurs because the multicast traffic is still on the old rpf interface, so a multicast discard route is installed. [PR1130238](#)
- IS-IS might flap during Routing Engine switchover. [PR1163770](#)

- See Also**
- [New and Changed Features on page 331](#)
 - [Changes in Behavior and Syntax on page 340](#)
 - [Known Behavior on page 343](#)
 - [Resolved Issues on page 346](#)
 - [Documentation Updates on page 360](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 361](#)
 - [Product Compatibility on page 365](#)

Resolved Issues

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

- [Resolved Issues: 15.1R6 on page 346](#)
- [Resolved Issues: 15.1R5 on page 348](#)
- [Resolved Issues: 15.1R4 on page 351](#)
- [Resolved Issues: 15.1R3 on page 354](#)
- [Resolved Issues: 15.1R2 on page 357](#)

Resolved Issues: 15.1R6

- [Class of Service \(CoS\) on page 347](#)
- [General Routing on page 347](#)
- [Infrastructure on page 347](#)
- [Multiprotocol Label Switching \(MPLS\) on page 347](#)
- [Platform and Infrastructure on page 348](#)
- [VPNs on page 348](#)

Class of Service (CoS)

- The following error log message might be seen with hierarchical CoS and strict-high scheduling configured. **Dec 27 11:08:02.293 mand-re0 fpc1 cos_check_temporal_buffer_status: IFD ge-1/2/1 IFL 358: Delay buffer computation incorrect.^M**. If hierarchical scheduler is configured for a physical interface and if guaranteed rate is not set for a logical interface under this physical interface, then the temporal buffer is configured. The display of error message is valid when guaranteed rate is 0, but it is not valid when guaranteed rate is disabled. [PR1238719](#)

General Routing

- After primary clock goes down, clock locks to previous secondary as new primary. The clock should lock to secondary. [PR1114281](#)
- On the FPC-SFF-PTX-P1-A (PTX3000), FPC-SFF-PTX-T (PTX3000), FPC-PTX-P1-A (PTX5000), and FPC2-PTX-P1A (PTX5000), packet loss might be observed in an equal-cost multipath (ECMP) or aggregated Ethernet scenario. It occurs in a race condition: because the unilist is created before ARP has learned MAC addresses, the selector table is corrupted. [PR1120370](#)
- In rare cases, multiple Routing Engine switchovers might result in a SNGPMB crash. (The SNG PMB is the same thing as SPMB. It is on the line card and contains the LCPU.) It also manages locally discovered issues and the switch fabric (through the chassis manager thread, which communicates with the fabric manager thread in chassisd). [PR1176094](#)
- For PTX Series routers, the IPv6 unilist next-hop member is in "replaced" status on the Packet Forwarding Engine after interface flapping with IPv6 Neighbor Discovery timeout. While the problem is occurring, the routing table will display all right next-hop status but cannot forward traffic because forwarding next-hop in Packet Forwarding Engine is in "replaced" status and no longer active. [PR1177023](#)
- When an ARP entry is learned through an AE interface and a route is pointing to that ARP nexthop, the ARP entry will not expire even if the ARP IP is not reachable. This issue occurs due to the route nexthop on the AE interface getting stuck in a unicast state even if the remote end is not reachable, and the RPD is unaware that the ARP is invalid. So, with this resolution, the route nexthop on the AE interface can be shown in the hold state when the remote end is not reachable. [PR1211757](#)

Infrastructure

- In an RSVP scenario, provision RSVP label-switched path (LSP) with ldp-tunneling enabled and these LSPs are configured with link protection, continuous kernel logs, and LDP stats timeout error might be seen when executing **show ldp traffic-statistics**. [PR1215452](#)

Multiprotocol Label Switching (MPLS)

- On a P2MP LSP transit router with link-protection enabled, if the LSP is the last sub-LSP, tearing the last sub-LSP (for example, a RESV tear message is received from a downstream router) might crash the routing process (rpd). [PR1036452](#)

- When we have a statically configured ingress and transit LSPs, due to timing issue, there could be a scenario wherein the selfID used by the transit LSP gets allocated to the ingress LSP. Ingress static LSP does not reuse the same selfID across restarts whereas transit static LSP tries to reuse the same. This leads to the RPD crash from the collision when the transit LSP tried to reuse the same ID. [PR1084736](#)
- On PTX Series platforms, the rpd might crash when the RSVP bypass undergoes re-optimization and the re-optimized instance encounters failure before it becomes the main instance. The core files can be seen by executing the CLI command **show system core files**. Stack trace: #0 0x0000000802ad8bd4 in patricia_node_in_tree () #0 0x0000000802ad8bd4 in patricia_node_in_tree () #1 0x00000000009ec3da in tag_pvc_shortwait () #2 0x0000000000a2fe94 in ted_delete_cc_from_link () #3 0x0000000000a3009d in ted2cspf_cleanup () #4 0x0000000000f27d56 in task_job_create_foreground () #5 0x0000000000f289e5 in task_job_bg_dispatch () #6 0x0000000000f24d85 in task_scheduler () #7 0x000000000062b9e2 in main (). [PR1250253](#)

Platform and Infrastructure

- There is a race condition between database creation and database access. This is rarely reproducible. There is no functional impact on the core. [PR1225086](#)

VPNs

- On PTX Series platform, the L2 circuit does not switch from primary to backup and vice versa based on the APS status change, because when APS switchover happens, the PW switchover does not switch to the new APS active neighbor. [PR1239381](#)

Resolved Issues: 15.1R5

- [General Routing on page 348](#)
- [Forwarding and Sampling on page 350](#)
- [Infrastructure on page 350](#)
- [MPLS on page 350](#)
- [Platform and Infrastructure on page 350](#)
- [Routing Protocols on page 351](#)
- [User Interface and Configuration on page 351](#)

General Routing

- To lock to the secondary node when primary node goes down, you should not reprogram the Centralized Clock Generator (CCG). However, when you are determining whether clock_selection should be aborted, if the old primary clock source has been removed from the configuration, do not abort; new sources need to be reselected. [PR1094106](#)
- The routing protocol process (rpd) fails to respond to any new CLI routing commands (for example, **show mpls lsp terse**). The rpd is forking a child process while rpd is processing a **show** command. When the subprocess tries to exit, it tries to close the management socket being used by the **show** command. This failure might cause the rpd subprocess to crash and generate a core file. It also removes the rpd pid file, which

prevents the rpd from processing any new CLI commands even though original rpd process continues to run normally. [PR1111526](#)

- On the FPC-SFF-PTX-P1-A(PTX3000), FPC-SFF-PTX-T(PTX3000), FPC-PTX-P1-A(PTX5000), and FPC2-PTX-P1A(PTX5000), packet loss might be observed in an equal-cost multipath (ECMP) or aggregated Ethernet (AE) scenario. It occurs in a race condition: because the unilist is created before ARP has learned MAC addresses, the selector table is corrupted. [PR1120370](#)
- On PTX series platforms with FPC3, the octets of IPv4 source and destination addresses in the firewall log are listed in reverse; this might affect troubleshooting. The IPv6 log works fine. This is a minor issue, and there is no other service impact. [PR1141495](#)
- Because of incorrect implementation in the code, power consumption was not fetched properly for the SIBs when using PTX PDU2. [PR1156265](#)
- FPC might crash after FPC reloading (restart FPC/non-GRES Routing Engine switchover), because of memory corruption when interface-specific filter process IPC messages. To fix this, the way firewall daemon (dfwd) for interface-specific filters is enhanced. Now, when the TLV decode has errors, the process discards the incorrectly decoded IPC message. [PR1164055](#)
- On PTX Series platforms, when a high-priority clock source (bits-a) goes down, the clock status transits from "locked to bits-a" to "holdover" to "acquiring" to "locked to bits-b". When the bits-a comes up, the clock status reverts from "locked to bits-b" to "holdover" to "acquiring" to "locked to bits-a". [PR1168000](#)
- For PTX Series routers, the IPv6 unilist next-hop member will become "replaced" status on Packet Forwarding Engine (PFE) after interface flapping with IPv6 ND (Neighbor Discovery) timeout. While the problem is happening, routing-table will display all right next-hop status but cannot forward traffic since forwarding next-hop in PFE is in "replaced" status and no longer active. [PR1177023](#)
- FPC might generate a core file when issuing clear threads and show threads simultaneously. [PR1184113](#)
- By default SNMP will cache SNMP values for 5 seconds. Sometimes the kernel will cache these values for a longer duration. [PR1188116](#)
- On PTX Series routers with FPC type 1 and FPC type 2, if there is a problem with ASIC in the FPC, the FPC might be disconnected from the Routing Engine. [PR1207153](#)
- In some conditions where the fan tray is not properly seated in PTX Series routers, the present PIN from the fan tray might not be detected and the fan tray is declared "Absent" in the output for the **show chassis environment** command. However, the alarm for this condition is not raised under "show chassis alarms" if the alarm occurs during a system reboot. [PR1216335](#)

Forwarding and Sampling

- The Sampling Route-Record Daemon (SRRD) process does not delete routes when the DELETE is received from RPD. This results in build-up of memory in SRRD daemon and once SRRD reaches the limit, it crashes and restarts itself. This scenario only occurs when one family is not configured on all the FPC clients (for example, FPC with inline J-Flow enabled or PIC with PIC-based sampling enabled in one client). Only IPv4 family is configured in all the clients, and IPv6 and MPLS families are not configured for sampling in any of the clients. [PR1180158](#)

Infrastructure

- When the kernel tries to collect statistics from a faulty FPC, it might trigger a kernel panic because of an invalid response from the faulty FPC. [PR1185013](#)

MPLS

- In the following scenario where 1) The PHOP link goes down and the router becomes MP for a LSP. 2) After some time, NHOP link for the same LSP goes down. The router becomes PLR for the same LSP. So effectively, the router is both MP and PLR for the same LSP. In this scenario, the router sends incorrect PathErr message for the backup MP PSB. It sends "Bad strict route" PathErr instead of "Tunnel local repaired" PathErr. [PR1132641](#)
- Changing the configuration under both [**protocols pcep**] and [**protocols mpls lsp-external-controller**] might trigger rpd to crash because of a race condition. [PR1194068](#)

Platform and Infrastructure

- When you configure one group with a configuration of routing-instances and apply that group under routing-instances, the rpd process crashes after executing you run the **activating routing-instances** or **deactivating routing-instances** commands. [PR1109924](#)
- In a very rare scenario, during TAC accounting configuration change, the auditd daemon crashes because of a race condition between auditd and its sigalarm handler. [PR1191527](#)

Routing Protocols

- A PTX Series node with a PR 1169289 fix might not be able to play the role of 6PE ingress node for inet6 traffic, if multipath is enabled for the peer giving the inet6 routes in the "inet6 labeled-unicast" family. This problem occurs because PR 1169289 causes the PTX Series router to create a composite next hop for the inet6.0 route, which is not supported. [PR1185362](#)

User Interface and Configuration

- When persist-groups-inheritance is configured and you issue a rollback, the configuration is not propagated properly after a commit. [PR1214743](#)

Resolved Issues: 15.1R4

- [Class of Service \(CoS\) on page 351](#)
- [General Routing on page 351](#)
- [High Availability \(HA\) and Resiliency on page 353](#)
- [Interfaces and Chassis on page 353](#)
- [MPLS on page 353](#)
- [Network Management and Monitoring on page 353](#)
- [Platform and Infrastructure on page 353](#)
- [Routing Protocols on page 353](#)
- [Software Installation and Upgrade on page 354](#)
- [VPNs on page 354](#)

Class of Service (CoS)

- In case of member links of an aggregated Ethernet (AE) interface scatter over multiple Packet Forwarding Engines, if the FPC where member links of the AE interface reside gets reset or the interface is disabled, there may be a dip in the output of SNMP walk on an AE-related queue MIB (such as jnxCosQstatTxedPkts). The behavior is intermittent and not seen every time. [PR1122343](#)
- This PR does optimization in AE SNMP handling. If all the links in an AE bundle go down, then any CoS SNMP query for this AE IFD/IFL will return cached values. [PR1140440](#)

General Routing

- FFP is a generic process that will be called during commit process, and FFP calls the PDB initialization as part of its process. On the PDB-unsupported platforms (MX Series, EX9200, M10i, M120, M320 is PDB-supported), when committing configuration, some error messages will be seen. [PR1103035](#)
- When a labeled BGP route resolves over a route with MPLS label (for example, LDP/RSVP routes), after clearing the LDP/RSVP routes, in the short window before the LDP/RSVP routes restore, if the BGP routes resolves over a direct route (for example, a one-hop LSP), the rpd process might crash. [PR1063796](#)

- When a switchover is done from one Routing Engine to the other, in graceful-switchover redundancy mode, there is a brief period early in the transition of the SIB to online state, during which unsolicited (not corresponding to an attempt by the CPU to access the SIB via PCIe) errors are received at the downstream PCIe port on the CB to the SIB. The fix is to mute the generation of such errors during this brief period of the switchover. [PR1068237](#)
- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)
- On PTX Series platforms with external clock synchronization interface configured, when both BITS external clocks are disconnected at the same time, the 100GbE-LR4 FINISAR interface might flap. This link flap issue is narrowed down to the operation of datapath FIFO within CFP. When both the BITS clocks are disconnected, the reference clock jumps to "free-running" mode. This transition leads to a phase shift in the reference clock. Because of this phase shift, the data rates into and out of the FIFO will temporarily not match, leading to a FIFO over-run or under-run condition. This over-run or under-run condition forces a FIFO reset, and the output signal is distorted. So the far-end interface detects a "local-fault," then returns a "remote-fault" back to the near-end, hence a link flap. Users need to manually configure the FPC recovered clock port for each clock put into "chassis synchronization source". Only one clock of each FPC can be put into "chassis synchronization source". [PR1091228](#)
- On PTX Series platforms, if there are scaling configurations (for example, 5,000 routes and each of them with 64 ECMP paths configured) on a single interface and an L2 rewrite profile is applied for the interface, the FPC might crash when deactivating and then activating the CoS configuration of the interface. [PR1096958](#)
- Entropy Label Capability is enabled by-default on all Juniper Networks (PTX Series and MX Series) systems. On PTX Series routers transit LSRs that carry LSPs with Entropy Label Capability, packet loss can be observed caused by data errors when one or more labeled route entries are not properly removed from the hash table (That is, following LSP optimization or MBB event) because the "stale" entries are pointing to corrupted route memory. As a result, when the MPLS label that is associated with the stale entry is reused, data errors are seen for packets using the corresponding label. [PR1100637](#)
- Because of a buffer size issue for FPC-SFF-PTX-P1-A (PTX3000) and FPC2-PTX-P1A (PTX5000), the "ISSU RECONNECT TIMEOUT" or "READY Message Without Reconnect" message is seen during unified ISSU. [PR1155936](#)

High Availability (HA) and Resiliency

- On MX Series platforms with Junos OS Release 15.1R1 or later, while a core file is being generated, if you try to access the dump file directory, the system might hang and crash due to the deadlock defect. [PR1087082](#)

Interfaces and Chassis

- During subscriber login or logout, the following error log might occur on the device configured with GRES/NSR: /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV. (pp0.1073751222) [PR1058958](#)

MPLS

- When an LSP is link-protected and has no-local-reversion configured, if the primary link (link1) is down and LSP is on bypass (link2), then another link (link3) is brought up, before the LSP switch to link3. If link1 is enabled and link3 is disabled, the LSP will remain in bypass LSP forever. This is a timing issue. [PR1091774](#)
- When multipoint LDP (M-LDP) in-band signaling is enabled to carry multicast traffic across an existing IP/MPLS backbone and routing process is enabled to use 64-bit mode, the rpd might crash because it is attempting to access an uninitialized local variables. [PR1118459](#)

Network Management and Monitoring

- While the router is rebooting and SNMP polling is not stopped, SNMP requests might land on the mib2d process before Routing Engine protocol mastership is resolved, causing the mib2d process crash. [PR1114001](#)

Platform and Infrastructure

- With the delta-export command enabled, "show|compare" output still appears after the last successful commit. [PR1129577](#)

Routing Protocols

- In an IS-IS environment MPLS LSPs are established, when the IS-IS traceoptions flag "general" is activated, and the LSP convergence time is increased. [PR1090752](#)
- In a multicast environment, when the rendezvous point (RP) is a first-hop router (FHR) with MSDP peers, when the rpf interface on the RP is changed to an MSDP-facing interface, traffic loss is seen. The loss occurs because the multicast traffic is still on the old rpf interface, so a multicast discard route is installed. [PR1130238](#)

Software Installation and Upgrade

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

VPNs

- For a Layer 2 circuit, the PTX3000 uses a different Virtual Circuit Connectivity Verification (VCCV) BFD control packet format from that of MX Series and the other PTX Series platforms. PTX3000 negotiates the router-alert control channel type and uses the PW Associated Channel Header of Channel Type : 0x0021. However, MX Series and the other PTX Series platforms use the channel Type 0x0007 without IP/UDP headers. Junos OS takes the Channel-type 0x0007 as default. MX Series and the other PTX Series platforms work as expected. This is a PTX3000-specific issue. [PR1116356](#)

Resolved Issues: 15.1R3

- [Class of Service \(CoS\) on page 354](#)
- [General Routing on page 355](#)
- [High Availability \(HA\) and Resiliency on page 356](#)
- [Interfaces and Chassis on page 356](#)
- [MPLS on page 356](#)
- [Network Management and Monitoring on page 356](#)
- [Platform and Infrastructure on page 356](#)
- [Routing Protocols on page 356](#)
- [Software Installation and Upgrade on page 357](#)
- [VPNs on page 357](#)

Class of Service (CoS)

- In case of member links of an aggregated Ethernet (AE) interface scatter over multiple Packet Forwarding Engines, if the FPC where member links of the AE interface reside get reset or the interface is disabled, there might be a decrease in the output of SNMP walk on the AE-related queue MIB (such as jnxCosQstatTxedPkts). The behavior is intermittent and not seen every time. [PR1122343](#)
- This PR does optimization in AE SNMP handling. If all the links in an AE bundle go down, then any CoS SNMP query for this AE IFD/IFL will return cached values. [PR1140440](#)

General Routing

- When a labeled BGP route resolves over a route with an MPLS label (for example, LDP/RSVP routes), after clearing the LDP/RSVP routes, in the short time before the LDP/RSVP routes restore, if the BGP routes resolve over a direct route (for example, a one-hop LSP), the rpd process might crash. [PR1063796](#)
- When a switchover is done from one Routing Engine to the other, in graceful-switchover redundancy mode, there is a brief period early in the transition of the SIB to online state, during which unsolicited (not corresponding to an attempt by the CPU to access the SIB through PCIe) errors are received at the downstream PCIe port on the CB to the SIB. The fix is to mute the generation of such errors during this brief period of the switchover. [PR1068237](#)
- On PTX Series platforms with external clock synchronization interface configured, when both BITS external clocks are disconnected at the same time, the 100GbE-LR4 FINISAR interface might flap. This link flap issue is narrowed down to the operation of data-path FIFO within CFP. When both the BITS clocks are disconnected, the reference clock jumps to "free-running" mode. This transition leads to a phase shift in the reference clock. Due to this phase shift, the data rates into and out of the FIFO will temporarily not match, leading to a FIFO over-run or under-run condition. This over-run or under-run condition forces a FIFO reset, and the output signal is distorted. So the far-end interface detects 'local-fault', then return 'remote-fault' back to the near-end, hence a link flap. User needs to manually configure FPC recovered clock port for each clock put into "chassis synchronization source". Only one clock of each FPC can be put into "chassis synchronization source". [PR1091228](#)
- On PTX Series platform, if there are scaling configurations (for example, 5000 routes and each of them with 64 ECMP paths configured) on a single interface and L2 rewrite profile is applied for the interface, the FPC may crash when deactivating and then activating the CoS configuration of the interface. [PR1096958](#)
- Starting with Junos Release 14.1, Entropy Label Capability is enabled by-default on all Juniper [PTX] systems. On PTX transit LSRs that carry LSPs with Entropy Label Capability, packet loss can be observed due to data errors when one or more labeled route entries are not properly removed from the hash table (i.e., following LSP optimization or MBB event) because the 'stale' entries are pointing to corrupted route memory. As a result, when the MPLS label that's associated with the 'stale' entry is re-used, data errors are seen for packets using the corresponding label. [PR1100637](#)
- FFP is a generic process that shall be called during commit process, and FFP calls the PDB initialization as part of its process. On the PDB-unsupported platforms (MX Series, EX9200, M10i, M120, M320 is PDB-supported), when committing configuration, some error messages will be seen. [PR1103035](#)

High Availability (HA) and Resiliency

- On PTX Series platform with Junos OS Release 15.1R1 and later, while a core dump is in progress, if we try to access the dump directory, due to the deadlock defect, the system might hang and crash. As a workaround, we should not access the "/var/crash" directory till the core dump is complete. [PR1087082](#)

Interfaces and Chassis

- During subscriber login/logout, the below error log might occur on the device configured with GRES/NSR. /kernel: if_process_obj_index: Zero length TLV! /kernel: if_pfe: Zero length TLV (pp0.1073751222). [PR1058958](#)

MPLS

- When an LSP is link-protected and has no-local-reversion configured, if the primary link (link1) is down and LSP on bypass (link2), then another link (link3) is brought up, before the LSP switch to link3, if link1 is enabled and link3 is disabled, the LSP will stuck in bypass LSP forever. This is a timing issue. [PR1091774](#)
- When multipoint LDP (M-LDP) in-band signaling is enabled to carry multicast traffic across an existing IP/MPLS backbone and routing process is enabled to use 64-bit mode, the rpd might crash due to accessing uninitialized local variables. [PR1118459](#)

Network Management and Monitoring

- While the router is rebooting and SNMP polling is not stopped, SNMP requests might land on mib2d process before Routing Engine protocol mastership is resolved, causing the mib2d process crash. [PR1114001](#)

Platform and Infrastructure

- The MIB counter or "show pfe statistics traffic" shows junk PPS and invalid total traffic output counter. [PR1084515](#)
- With delta-export command enabled, "show |compare" output still shows after last successful commit. [PR1129577](#)

Routing Protocols

- In IS-IS environment, MPLS LSPs are established, when IS-IS traceoptions flag "general" is activated, the LSP convergence time is increased. [PR1090752](#)
- In multicast environment, when the RP is FHR (first-hop router) and it has MSDP peers, when the rpf interface on RP changed to MSDP facing interface, due to the multicast traffic is still on the old rpf interface, a multicast discard route will be installed and traffic loss will be seen. [PR1130238](#)

Software Installation and Upgrade

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

VPNs

- For Layer 2 circuit, PTX3000 uses different VCCV (Virtual Circuit Connectivity Verification) BFD control packet format from that of MX Series and the other PTX Series platforms. PTX3000 negotiates Router-alert control channel type, and uses PW Associated Channel Header of Channel Type : 0x0021. However, MX Series and the other PTX Series platforms use the Channel Type is 0x0007 without IP/UDP headers. Junos OS takes the Channel-type 0x0007 as default. MX Series and the other PTX Series platforms work as expected. This is a PTX3000 specific issue. [PR1116356](#)

Resolved Issues: 15.1R2

- [Forwarding and Sampling on page 357](#)
- [General Routing on page 357](#)
- [Interfaces and Chassis on page 358](#)
- [MPLS on page 359](#)
- [Network Management and Monitoring on page 359](#)
- [Routing Protocols on page 359](#)

Forwarding and Sampling

- In PTX Series Carrier-Grade Service Engine (CSE) jflow solution environment, because the sampling process (sampled) may get into a continuous loop when handling asynchronous event (for example, aggregated tethered services interface flapping, or route update, or IFL/IFD update), the sampled may never come out of that loop which may result in high CPU usage (up to 90 % sometimes). Because, sampled is not able to consume any states (such as route updates, interface updates) generated by kernel and this results in memory exhaustion, finally resulting in the router not making any updates and forcing a router reboot. [PR1092684](#)

General Routing

- On PTX Series platform, when performing scaling (for example, polling 768 IFDs via SNMP with max of 92 PPS and with all 8 FPCs online) SNMP polling on the device, due to the large number of messages between Routing Engine and Packet Forwarding Engine, PFEMAN (Packet Forwarding Engine manager) errors might be seen on the router, which may cause high SNMP response time and CPU spike (for example, increase 8 % when executing the "show" command) as well. [PR1078003](#)
- On PTX3000 routers running Junos OS Release 14.1 and later, the Packet Forwarding Engine does not support L3VPN VRF. For example, when you assign the loopback (lo0) interface to VRF as the management VRF, the following commit error is returned: **# commit check [edit routing-instances l3vpn interface] 'et-8/0/0.0' RT Instance: Only loopback interface is supported under vrf routing instances. error: configuration check-out**

failed Note that in Junos OS Release 14.2, you will see the same commit error, but the commit will be successful. You might also encounter a packet discard issue. [PR1078960](#)

- Tunable SFP+ optics will not be supported on P1-PTX-24-10G-W-SFPP PIC in Junos OS 15.1R1 release. On Tunable Optics in this PIC, with 15.1R1, the wavelength will not be configurable and the tunable parameters will not be correctly displayed in the CLI. [PR1081992](#)
- The FPC on PTX Series router might crash and reboot when the Packet Forwarding Engine is handling a fatal error; when the error happened, "TQCHIP0: Fatal error pqt_min_free_cnt is zero" log message will be seen. [PR1084259](#)
- Due to a change in an existing PR, group names in the configuration must be a string of alphanumericals, dashes, or underscores. There is not workaround other than following the group name instructions. [PR1087051](#)
- In Dual Routing Engine systems when both Routing Engines reboot and after coming up, if the mastership is not established or takes time to establish, mib2d may start and exit four times in quick succession. Hence it will not be running. As a workaround, it can be simply started again once Routing Engine mastership is established. This is a race condition and hence may not be seen always. [PR1087428](#)
- On PTX Series platforms, some non-fatal interrupts (for example, CM cache or AQD interrupts) are logged as fatal interrupts. The following log messages will be shown on CM parity interrupt: fpc0 TQCHIP 0: CM parity Fatal interrupt, Interrupt status: 0x10 fpc0 CMSNG: Fatal ASIC error, chip TQ fpc0 TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180010 msec TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180005 msec [PR1089955](#)
- On Junos OS Release 15.1R1, when the multicast next-hop is changed, the grafting and pruning operations take more time than before. [PR1090608](#)
- When the PTX Series only has bits-a and bits-b as configured clock sources (and there is no interface on FPC configured as clock source), and it is losing signal from both of bits-a and bits-b simultaneously, clock sync state will go to FREERUN mode immediately, this is unexpected behavior. After the fix of this PR, clock sync state will stay HOLDOVER, then will go to FREERUN mode after the timeout. [PR1099516](#)
- On PTX Series platform, when yanking out FPC or SIB ungracefully (for example, pulling the line card out of the chassis unintentionally when the line card is carrying the traffic), there might be small probability that it can impact any of the FPCs with Grant Scheduler (GS) and Request Table (RT) fatal interrupt occurred. [PR1105079](#)

Interfaces and Chassis

- If we load the 15.1 Junos jinstall/jinstall64 image on PTX Series and if we have CFM configured over AE interfaces, the FPC might crash. [PR1085952](#)
- In the dual Routing Engines scenario with GRES and ae0 interfaces configuration, if GRES is disabled on system, the backup Routing Engine should remove the ae0 bundle; however, it does not go clean and ae0 remains in the backup Routing Engine. After switching Routing Engine mastership to make other Routing Engine as master, the new master Routing Engine (which was backup earlier) continues to use invalid MAC address "00:00:00:00:00:00". [PR1089946](#)

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

MPLS

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

Network Management and Monitoring

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

Routing Protocols

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

- See Also**
- [New and Changed Features on page 331](#)
 - [Changes in Behavior and Syntax on page 340](#)
 - [Known Behavior on page 343](#)
 - [Documentation Updates on page 360](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 361](#)
 - [Product Compatibility on page 365](#)

Documentation Updates

This section lists the errata and changes in Junos OS Release 15.1R6 documentation for the PTX Series.

- [High Availability Feature Guide on page 360](#)
- [IPv6 Neighbor Discovery Feature Guide on page 360](#)

High Availability Feature Guide

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

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

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

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

IPv6 Neighbor Discovery Feature Guide

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

- See Also**
- [New and Changed Features on page 331](#)
 - [Changes in Behavior and Syntax on page 340](#)
 - [Known Behavior on page 343](#)
 - [Known Issues on page 343](#)
 - [Resolved Issues on page 346](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 361](#)
 - [Product Compatibility on page 365](#)

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the PTX Series. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

- [Upgrading Using Unified ISSU on page 361](#)
- [Upgrading a Router with Redundant Routing Engines on page 361](#)
- [Basic Procedure for Upgrading to Release 15.1 on page 361](#)

Upgrading Using Unified ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [High Availability Feature Guide for Routing Devices](#).

For information on ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\) Web application](#).

Upgrading a Router with Redundant Routing Engines

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

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

For the detailed procedure, see the [Installation and Upgrade Guide](#).

Basic Procedure for Upgrading to Release 15.1

When upgrading or downgrading Junos OS, use the **jinstall** package. For information about the contents of the **jinstall** package and details of the installation process, see the [Installation and Upgrade Guide](#). Use other packages, such as the **jbundle** package, only when so instructed by a Juniper Networks support representative.



.....

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

.....



.....

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

.....



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

The download and installation process for Junos OS Release 15.1R6 is different from previous Junos OS releases.

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



NOTE: After you install a Junos OS Release 15.1R6 jinstall package, you cannot issue the request system software rollback command to return to the previously installed software. Instead you must issue the request system software add validate command and specify the jinstall package that corresponds to the previously installed software.

The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is

a different release. Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes. Rebooting occurs only if the upgrade is successful.

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

```
user@host> request system software add validate reboot source/jinstall-15.1
R61-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-15.1
R61-export-signed.tgz
```

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

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

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

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

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 15.1 jinstall package, you cannot issue the `request system software rollback` command to return to the previously installed software. Instead you must issue the `request system software add validate` command and specify the jinstall package that corresponds to the previously installed software.

- See Also**
- [New and Changed Features on page 331](#)
 - [Changes in Behavior and Syntax on page 340](#)
 - [Known Behavior on page 343](#)
 - [Known Issues on page 343](#)

- [Resolved Issues on page 346](#)
- [Documentation Updates on page 360](#)
- [Product Compatibility on page 365](#)

Product Compatibility

- [Hardware Compatibility on page 365](#)

Hardware Compatibility

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

To determine the features supported on PTX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at:
<https://pathfinder.juniper.net/feature-explorer/>

- See Also**
- [New and Changed Features on page 331](#)
 - [Changes in Behavior and Syntax on page 340](#)
 - [Known Behavior on page 343](#)
 - [Known Issues on page 343](#)
 - [Resolved Issues on page 346](#)
 - [Documentation Updates on page 360](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 361](#)

Junos OS Release Notes for the QFX Series

These release notes accompany Junos OS Release 15.1R6 for the QFX 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 https://www.juniper.net/documentation/product/en_US/junos-os.

- [New and Changed Features on page 366](#)
- [Changes in Behavior and Syntax on page 369](#)
- [Known Behavior on page 370](#)
- [Known Issues on page 374](#)
- [Resolved Issues on page 375](#)
- [Documentation Updates on page 390](#)
- [Migration, Upgrade, and Downgrade Instructions on page 390](#)
- [Product Compatibility on page 394](#)

New and Changed Features

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



NOTE: The following QFX Series platforms are supported in Release 15.1R6: QFX3500, QFX3600, and QFX5100.

- [Management on page 366](#)
- [Network Management and Monitoring on page 368](#)
- [Spanning-Tree Protocols on page 368](#)
- [User Interface and Configuration on page 368](#)

Management

- **Support for YANG features including configuration hierarchy `must` statement constraints published in YANG, and a module that defines Junos OS YANG extensions (QFX Series)**—Starting with Junos OS Release 15.1R3, 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 Juniper Networks 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 (QFX Series)**—Starting with Junos OS Release 15.1R3, 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.](#)]

Network Management and Monitoring

- **Monitor Virtual Chassis ports (VCPs) with SNMP (QFX3500, QFX3600)**—Starting with Junos OS Release 15.1R3, you can configure the switch to monitor VCPs with SNMP. To enable SNMP monitoring of VCPs in a Virtual Chassis or Virtual Chassis Fabric (VCF), use the **set virtual-chassis vcp-snmp-statistics** CLI command.

Spanning-Tree Protocols

- **Global configuration of spanning-tree protocols (QFX Series)**—Starting with Junos OS Release 15.1R13, global configuration of the spanning-tree protocols RSTP, MSTP, and VSTP is supported on QFX 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](#).]

User Interface and Configuration

- **Support for replacing patterns in configuration data within NETCONF and Junos OS XML protocol sessions (QFX Series)**—Starting with Junos OS Release 15.1R3, you can replace variables and identifiers in the candidate configuration when performing a **<load-configuration>** operation in a Junos OS 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 Also**
- [Changes in Behavior and Syntax on page 369](#)
 - [Known Behavior on page 370](#)
 - [Known Issues on page 374](#)
 - [Resolved Issues on page 375](#)
 - [Documentation Updates on page 390](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 390](#)
 - [Product Compatibility on page 394](#)

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 15.1R6 for the QFX Series.

- [Interfaces and Chassis on page 369](#)
- [Routing Protocols on page 370](#)

Interfaces and Chassis

- **Configuring unified forwarding table profiles (EX4600 Virtual Chassis, QFX5100 Virtual Chassis, and QFX Series Virtual Chassis Fabric)**—Starting in Junos OS Release 15.1R5, Packet Forwarding Engines on switches in a Virtual Chassis or Virtual Chassis Fabric (VCF) do not automatically restart upon configuring and committing a unified forwarding table profile change using the **set chassis forwarding-options** statement. Instead, a message is displayed at the CLI prompt and logged to the switch's system log, prompting you to reboot the Virtual Chassis or VCF for the change to take effect. This change avoids Virtual Chassis or VCF instability that might occur with these switches if the profile update propagates to member switches and otherwise causes multiple Packet Forwarding Engines to automatically restart at the same time. This behavior change does not apply to other switch types or to EX4600 and QFX5100 switches not in a Virtual Chassis or VCF; in those cases, the switch continues to restart automatically when a unified forwarding table profile change is committed.

We recommend that you plan to make profile changes in a Virtual Chassis or VCF comprised of these switches only when you can perform a Virtual Chassis or VCF system reboot shortly after committing the configuration update, to avoid instability if one or more member switches restart unexpectedly with the new configuration (while the remaining members are still running the old configuration).

[See [Configuring the Unified Routing Table](#) and [forwarding-options \(chassis\)](#).]

- **New vc-path command display for Virtual Chassis Fabric (VCF)**—Starting in Junos OS Release 15.1R5, the output from the **show virtual-chassis vc-path** command displays additional fields when showing the forwarding path from a source interface to a destination interface in a Virtual Chassis Fabric (VCF), including details of multiple possible next hops. The **vc-path** command display for a forwarding path in a Virtual Chassis remains unchanged.

[See [show virtual-chassis vc-path](#).]

Routing Protocols

- **Support for RFC 6996, RFC 7300, and Internet draft draft-ietf-idr-as0-06 (QFX Series)**—Starting with Junos OS Release 15.1, RFC 6996, *Autonomous System (AS) Reservation for Private Use*, RFC 7300, *Reservation of Last Autonomous System (AS) Numbers*, and Internet draft *draft-ietf-as0-06* are supported.

RFC 7300, *Reservation of Last Autonomous System (AS) Numbers*, and the Internet draft *draft-ietf-idr-as0-06* restrict the use of 2-byte AS number 65535, 4-byte AS number 4294967295UL, and AS number 0 in a configuration. When you use these restricted AS numbers, the commit operation fails.

- See Also**
- [New and Changed Features on page 366](#)
 - [Known Behavior on page 370](#)
 - [Known Issues on page 374](#)
 - [Resolved Issues on page 375](#)
 - [Documentation Updates on page 390](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 390](#)
 - [Product Compatibility on page 394](#)

Known Behavior

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 15.1R6 for the QFX 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](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Multicast Protocols](#)
- [Multiprotocol Label Switching \(MPLS\)](#)
- [Platform and Infrastructure](#)
- [Routing Policy](#)
- [Routing Protocols](#)
- [Software-Defined Networks \(SDN\)](#)
- [Software Installation and Upgrade](#)
- [Spanning-Tree Protocols](#)
- [Virtual Chassis and Virtual Chassis Fabric \(VCF\)](#)
- [VPNs](#)

High Availability (HA) and Resiliency

- On QFX5100 switches, Fibre Channel over Ethernet (FCoE) traffic might be dropped for up to 4 seconds during an in-service software upgrade (ISSU) when FCoE Initialization Protocol (FIP) snooping is enabled. [PR981306](#)
- On EX4600 and QFX5100 switches, the Link Aggregation Control Protocol (LACP) in either slow mode or fast mode might go down and then come back up, causing a timeout and a service outage during an in-service software upgrade (ISSU) or a nonstop software upgrade (NSSU). In addition, after the master Routing Engine is rebooted, the switches might experience intermittent traffic loss on non-LAG interfaces, and redundant trunk group (RTG) convergence times might be long. [PR1031338](#)

Interfaces and Chassis

- On an EX4300 or a QFX5100 switch, a MAC address that is specified as part of a MAC-based VLAN is authenticated on an interface (for example, on xe-1/1/1), on which 802.1X authentication in multiple supplicant mode is configured. However, the same MAC address might not be authenticated on another interface, for example, xe-2/1/1, if the MAC address moves to interface xe-2/1/1 from interface xe-1/1/1. [PR1007589](#)
- On an MC-LAG, if an ARP for a host is learned on the MC-LAG interface and the host changes its MAC address without sending a gratuitous ARP, traffic loss might occur. [PR1009591](#)
- On QFX5100 switches, if you configure MC-LAG, IRB mac sync, and LACP force up, the number of packets received (rx) might be twice the amount sent (tx) from the customer edge to the core. [PR1015655](#)
- On a QFX5100 switch, you might be unable to commit the configuration if you modify the subnet of an IP address on an IRB interface by using the **replace pattern** command. [PR1119713](#)
- On QFX5100 Virtual Chassis, generic routing encapsulation (GRE) counters might not increment with a firewall filter and PIM configured. [PR1124170](#)
- On a QFX5100 or EX4600 switch, high ICMP delays are experienced when pinging directly connected integrated routing and bridging (IRB) interfaces. This is caused by a hardware limitation. Transit traffic is not affected. [PR1164135](#)
- On QFX5100 switches, Layer 2 control frames with a destination MAC address of 01:80:c2:00:00:02 and an ethertype of 8809 might be dropped at the egress PE router Layer 2 VPN. [PR1182124](#)

Layer 2 Features

- On a mixed-mode Virtual Chassis Fabric (VCF) with interface-mac-limit configured, if you remove the complete mac-limit configuration, the mac-limit behavior might remain. As a workaround, try rebooting the device. [PR1044460](#)
- On ELS (Enhanced Layer 2 Software) platforms (including EX4300, EX4600, EX9200, QFX3500, QFX3600, and QFX5100 switches), if Q-in-Q tunneling is enabled, if you

configure an RTG (redundant trunk group) on a Q-in-Q interface, the RTG configuration cannot be applied; there is a commit check error. [PR1134126](#)

- On QFX5100 switches with a CoS classifier configured on an AE interface, if you add or delete a subinterface, traffic loss of approximately 10 packets might occur while you are committing the changes. [PR1162963](#)
- On a QFX5100 switch, with a fully meshed MC-LAG topology configured, sometimes there is more traffic loss when the ICL interface goes down and then back up compared with when you have Junos OS Release 14.1X53-D35 software installed. The root cause has been identified, and this issue does not affect MC-LAG functionality. [PR1209322](#)

Multicast Protocols

- When an IGMP leave is sent from a host to a QFX5100 switch, one packet per multicast group is dropped during route programming. [PR995331](#)

Multiprotocol Label Switching (MPLS)

- On a QFX5100 switch, if an MPLS link is in hot standby mode and a pseudowire switchover is triggered by the event "remote site local interface signaled down," traffic flowing through the pseudowire might drop. [PR1027755](#)

Platform and Infrastructure

- Traffic convergence delay time for link protection, node-link protection, and fast reroute is more than 50ms for the QFX5100-48T switch. [PR1026957](#)

Routing Policy

- On the QFX Series, in a BGP equal-cost multipath (ECMP) scenario, if the import policy uses the policy action **next-hop peer-address** to set the route's protocol next-hop, BGP multipath might use more ECMP groups than necessary. If the ECMP entries exceed the maximum supported by the hardware, traffic loss might occur. As a workaround, use the policy action **next-hop ip-address** instead of the action **next-hop peer-address**.

Routing Protocols

- On a QFX Series Virtual Chassis, if you delete a member of a LAG associated with an IRB interface, the counter for the filter applied to the IRB interface might reset. [PR898171](#)
- On EX4300, EX4600, and QFX Series switches, a Bidirectional Forwarding Detection (BFD) session might not come up when BFD version 0 is configured. As a workaround, deactivate or delete the version configuration. [PR1076052](#)

Software-Defined Networks (SDN)

- On QFX5100 switches, if more than 1K virtual extensible LAN network identifiers (VNIs) are created by Open vSwitch Database (OVSDB), the VTEP gateway daemon (vgd) might generate a core file. [PR1075189](#)

Software Installation and Upgrade

- On EX4600, QFX3500, and QFX5100 switches, the amount of time that it takes for Zero Touch Provisioning to complete might be lengthy because TFTP might take a long time to fetch required data. [PR980530](#)
- On EX Series or QFX Series Virtual Chassis or Virtual Chassis Fabric (VCF), nonstop software upgrade (NSSU) cannot be used to upgrade from a Junos OS Release 14.1X53 image to a Junos OS Release 15.1 or later image. [PR1087893](#)

Spanning-Tree Protocols

- On QFX5100 Virtual Chassis interfaces on which the **flexible-vlan-tagging** statement is specified, STP, RSTP, MSTP, and VSTP are not supported. [PR1075230](#)

Virtual Chassis and Virtual Chassis Fabric (VCF)

- In a mixed Virtual Chassis or Virtual Chassis Fabric (VCF), the **show pfe filter hw summary** command is not supported for an EX4300 member of the Virtual Chassis or VCF. [PR1019377](#)
- On a QFX5100, QFX3600, QFX3500, or EX4300 switch, if you remove a transceiver from an interface and then reinsert it in the interface within 30 seconds after you have issued the **set virtual-chassis vc-port set** command to convert the interface into a Virtual Chassis port (VCP), the VCP is not created. [PR1029829](#)
- On a QFX5100 Virtual Chassis, frequent MAC move events can put the system into an inconsistent state, which results in a Packet Forwarding Engine manager (FXPC) process crash with a core file generated. [PR1086108](#)
- On QFX3500 and QFX3600 Virtual Chassis, any change in channelization causes the Packet Forwarding Engine to restart. If you apply channelization across various member switches in the Virtual Chassis, connectivity might be lost temporarily. [PR1105371](#)
- In a mixed mode Virtual Chassis with QFX3500 switches, if multicast packets are sent to the Routing Engine at a high rate, the Virtual Chassis might become unresponsive. [PR1117133](#)

VPNs

- On a QFX5100 switch that has performed a pseudowire switchover, traffic might drop for 10 seconds immediately after the switchover. [PR1049606](#)

- See Also**
- [New and Changed Features on page 366](#)
 - [Changes in Behavior and Syntax on page 369](#)

- [Known Issues on page 374](#)
- [Resolved Issues on page 375](#)
- [Documentation Updates on page 390](#)
- [Migration, Upgrade, and Downgrade Instructions on page 390](#)
- [Product Compatibility on page 394](#)

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 15.1R6 for the QFX Series.

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

- [Routing Protocols on page 374](#)
- [Security on page 374](#)
- [Spanning Tree Protocols on page 374](#)
- [Virtual Chassis and Virtual Chassis Fabric on page 374](#)

Routing Protocols

- On QFX5100, when resilient hashing is enabled on ECMP paths, flows on other paths should not be rehashed when one path goes down. But for host routes (/32 routes), rehashing might happen in some cases. [PR1137998](#)

Security

- On EX4300, EX4600, and QFX5100 switches, when a VLAN is mirrored, the mirrored packets may contain 38 additional bytes. The IP address in this packet is randomly generated and may appear as one of many existing, valid IP addresses on the Internet. It may appear as ERSPAN as well, which is a proprietary non-Juniper protocol. These addresses and packet types can be ignored. They may appear as alerts in certain IDP / IDS's and in packet analyzer applications, which you can ignore. [PR1170589](#)

Spanning Tree Protocols

- On QFX5100 Virtual Chassis interfaces on which the **flexible-vlan-tagging** statement is specified, STP, RSTP, MSTP, and VSTP are not supported. [PR1075230](#)

Virtual Chassis and Virtual Chassis Fabric

- On a Virtual Chassis Fabric, Virtual Chassis ports (VCPs) internal traffic looping causing traffic loss might be seen for known multicast traffic with TTL=1. [PR1042270](#)
- A VCF might not communicate properly with the backup spine when it has the configuration parameter **forwarding-options enhanced-hash-key fabric-load-balance flowlet** configured while upgrading. [PR1141965](#)

- See Also**
- [New and Changed Features on page 366](#)
 - [Changes in Behavior and Syntax on page 369](#)
 - [Known Behavior on page 370](#)
 - [Resolved Issues on page 375](#)
 - [Documentation Updates on page 390](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 390](#)
 - [Product Compatibility on page 394](#)

Resolved Issues

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

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

- [Resolved Issues: Release 15.1R6 on page 375](#)
- [Resolved Issues: Release 15.1R5 on page 377](#)
- [Resolved Issues: Release 15.1R4 on page 380](#)
- [Resolved Issues: Release 15.1R3 on page 383](#)

[Resolved Issues: Release 15.1R6](#)

- [High Availability \(HA\) and Resiliency](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Multiprotocol Label Switching \(MPLS\)](#)
- [Network Management and Monitoring](#)
- [Port Security](#)
- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Software-Defined Networking \(SDN\)](#)

High Availability (HA) and Resiliency

- On QFX5100 and EX4600 switches, during a nonstop software upgrade (NSSU), if an aggregated Ethernet (AE) interface is configured with multiple subinterfaces across multiple Flexible PIC Concentrators (FPCs), the AE interface might go down. [PR1227522](#)
- On a QFX5100 switch, you cannot perform an in-service software upgrade from Junos OS Release 14.1X53-D30 to Junos OS Release 14.1X53-D40. As a workaround, during a maintenance window, download the new software version, perform a regular software upgrade, and reboot the switch. [PR1229272](#)

Interfaces and Chassis

- On QFX5100 Virtual Chassis, DHCPv6 binding might fail if the server and the client are in different virtual routing and forwarding (VRF) instances. [PR1167693](#)
- Output from **show chassis environment** says fan tray testing/absent in QFX3500 Virtual Chassis with EX4300. [PR1200638](#)
- The backup link in the aggregated link is not forwarding the traffic when the primary link goes down in the following configuration with Junos OS Release 15.1R4: **root# show interfaces ge-0/0/10 ether-options { 802.3ad { ae0; primary; } } {master:0}[edit] root# show interfaces ge-0/0/19 ether-options { 802.3ad { ae0; backup; } }**. [PR1208614](#)
- On QFX Series switches, LLDP does not work on management and internal Ethernet (em) interfaces. [PR1224832](#)
- On QFX Series switches, in rare cases, the Link Up / Down notification from the Packet Forwarding Engine (PFE) to the Routing Engine might need a bit of time, so the PFE-side interface and remote device interface show Admin Up and Link Up, but the CLI might show the interface in Admin Down and Link Down. When this issue happens, it might last about 30 seconds. [PR1227947](#)
- A QFX5100-48S or QFX5100-96S might incorrectly list the media type of an SFP-T copper module as “fiber” in the output of **show interface**. [PR1240681](#)

Layer 2 Features

- On QFX5100 switches, if you configure an aggregated Ethernet (AE) interface in a VLAN associated with a VNI, the AE interface might stop forwarding traffic. Also, even after you delete the VXLAN configuration, the problem persists. [PR1213701](#)
- On QFX5100 switches, an fxpc process might generate a core file. [PR1231071](#)
- MAC learning will be very slow when clearing MAC addresses in cases of scale MAC learning (128k). [PR1240114](#)

Multiprotocol Label Switching (MPLS)

- Ping over LSP shows different behavior in regards to HLIM. [PR1179518](#)
- On EX Series and QFX Series switches, if you change a Layer 2 circuit configuration from Ethernet CCC encapsulation to VLAN CCC encapsulation, traffic losses might occur at the pseudowire tunnel initiation point. [PR1222888](#)

Network Management and Monitoring

- Despite the EX4300 switch or QFX5100 switch being configured with the network analytics feature, the analytics process might not run. As a result, the network analytics feature might be unable to collect traffic, queue statistics, and generate reports. [PR1165768](#), [PR1184720](#)
- The Digital Optical Monitoring (DOM) MIB jnxDomCurrentTable for 1G SFP interfaces does not return any value. [PR1218134](#)

Port Security

- On QFX3500, DHCP binding might not work when untrusted ARP inspection is enabled in the snooping device. [PR1229399](#)

Routing Policy and Firewall Filters

- On QFX5100 switches, firewall filters that contain policers might not process packets correctly if TCAM entries are programmed over multiple slices of TCAM memory space. Firewall filter terms are programmed as TCAM entries in the TCAM memory table. The auto-expansion function over multiple slices might fail with policers being attached to firewall filter terms. [PR1232926](#)

Routing Protocols

- In a QFX5100 Virtual Chassis or Virtual Chassis Fabric (VCF), if the master Routing Engine crashes when nonstop active routing (NSR) is configured and the **[edit system] switchover-on-routing-crash** statement is set, the Virtual Chassis or VCF fails to perform the switchover to the backup Routing Engine. The **switchover-on-routing-crash** statement helps to prevent loss of traffic during a Routing Engine switchover when NSR is enabled by switching immediately over to the backup Routing Engine. [PR1220811](#)
- On EX4600 and QFX Series switches with **unicast-in-lpm** configured, EBGp packets with ttl=1 and non-EBGP packets with ttl=1, whether destined for the device or even transit traffic, both go to the same queue. This might result in dropping of valid EBGp packets, resulting in EBGp flap. [PR1227314](#)
- On QFX5100 switches running Junos OS Release 14.1X53-D30.3, when you apply an IPv6 firewall filter, the system might crash with a PFE panic. [PR1234729](#)
- On a QFX5100 switch, Gratuitous Address Resolution Protocol (GARP) reply packets are not updating the Address Resolution Protocol (ARP) table. GARP request packets, however, are updating the ARP table as expected. [PR1246988](#)
- On QFX5100 switches, multicast route leaking does not support a Layer 3 interface (IPv4) as an upstream port. As a workaround, use an integrated routing and bridging (IRB) interface. [PR1250430](#)

Software-Defined Networking (SDN)

- On QFX5100 switches, OVSDb traffic might be dropped after Layer 2 learning is restarted. [PR1177012](#)

Resolved Issues: Release 15.1R5

- [Class of Service \(CoS\)](#)
- [Firewall Filters](#)
- [Infrastructure](#)
- [MPLS](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)

- [Spanning-Tree Protocols](#)
- [Virtual Chassis](#)

Class of Service (CoS)

- In an ETS configuration, if transmit-rate is configured at queue-level, the guaranteed rate should be configured at the TCP level. If not, a syslog message is logged about configuration failure. The configuration is not pushed to the kernel/PFE. On a QFX5100 Virtual Chassis, when a member joins, since the configuration check is already done on the master, the configuration is sent to members. Because the guaranteed rate is configured as 0, the logic to calculate the transmit-rate fails. [PR1195498](#)

Firewall Filters

- On QFX5100 switches, the DSCP action modifier of a **family inet** firewall filter does not properly modify or mark the DSCP bits on packets matching the firewall filter. [PR1205072](#)
- On QFX5100 switches, **port-range-optimize** (both source and destination) might fail to be programmed into the hardware for an inet output filter. [PR1211576](#)

Infrastructure

- On QFX5100 and EX4600 switches, in a rare timing condition, if there was already a request to gather some info from the QSFP and remove it at the same time, the Packet Forwarding Engine manager (fxpc) might crash. [PR1151295](#)
- On an EX4300 switch in a VCF, if a Layer 3 AE interface is looped back with a Layer 2 port in the same VLAN, then traffic with the same destination MAC to the AE interface is dropped (for example, the ping address of the AE interface). [PR1157283](#)
- On QFX5100-48T, when issuing **show interface extensive** or **show interface media**, the Local resolution: section of the Autonegotiation information section indicates that flow control is enabled for both tx and rx even though flow control has been explicitly configured as disabled and the disabled state is indicated in the top portion of the output. [PR1168511](#)
- On QFX5100 switches, packet loss and framing errors might be observed on QSFP+40GE-LX4 transceiver. [PR1177499](#)
- On EX4300, EX4600, QFX3500, QFX3600, and QFX5100 switches with **vlan-rewrite** configured on an AE interface, a VLAN rewrite might fail and result in traffic loss. [PR1186821](#)
- On QFX5100 switches that are running with VXLAN Open vSwitch Database (OVSDb), the Packet Forwarding Engine manager (fxpc) might crash and generate a core file because of heap memory exhaustion on the kernel. This is a specific issue with OVSDb and does not affect multicast VXLAN. [PR1187299](#)
- After you add or remove a PEM on a QFX5100 switch, the **show chassis environment pem** command does not display the correct Current(A) and Power(A) usage. [PR1204850](#)
- If a QFX5100 switch or VCF is configured with IGMP snooping without any PIM-related configuration, a mcsnoopd memory leak might occur when the device receives PIM hello packets that need to be forwarded further. When PIM hello packets are arriving

on the device, 12 bytes are allocated for every PIM hello packet, causing an increase in the memory consumed by the mcsnoopd process. [PR1209773](#)

MPLS

- On QFX5100 switches or a QFX3500 or QFX3600 Virtual Chassis, IP packet frames of 1500 bytes might drop when **family mpls** is configured on a logical interface. [PR1199919](#)
- On QFX5100 switches with MPLS and LDP enabled, for packets with incoming labels that must perform a penultimate hop popping (PHP) operation on the QFX5100 switch, occasionally the packets are not processed and are dropped. [PR1190437](#)

Platform and Infrastructure

- The Packet Forwarding Engine manager daemon (fxpc) might crash on an QFX5100 switch if multiple processes attempt to access the Ethernet-switching table/database at the same time. [PR1146937](#)
- On EX4600 or QFX5100 switches or Virtual Chassis or Virtual Chassis Fabric (VCF), when you reconfigure or modify the Unified Forwarding Table (UFT) profile, the device automatically restarts (for the UFT configuration to take effect). When this happens in a Virtual Chassis or Virtual Chassis Fabric (VCF) environment, the Virtual Chassis or VCF might become unstable and fail to recover, and the Virtual Chassis or VCF (all member devices) must be rebooted to reestablish stable operation. To avoid this situation, configure the UFT profile when you initially set up the device. After the fix, for standalone switches and Virtual Chassis with a single member, it works as before. For a Virtual Chassis or VCF with more than one member, the member does not restart, and the system generates a syslog message that tells you to restart the system manually when you change the UFT configuration. [PR1152102](#)
- On QFX3500 or QFX5100 switches, when parity errors occur on interfaces, they might affect the memory management unit (MMU) memories. MMU counters can be corrupted, the interface buffers might be stuck, and there might be interface flaps and traffic loss on the affected ports. As a workaround (restoration only), reboot the system. [PR1169700](#)
- In a QFX5100 Virtual Chassis, if the master is halted or rebooted with some limited MAC persistence timer set, then in a specific sequence the IRB MAC does not get programmed correctly in the BCM. [PR1188092](#)
- On QFX3500, QFX3600, QFX5100, and EX4600 switches, if a routing loop is created, the TTL of the packet does not reduce to 0 and the packet is not dropped. [PR1196354](#)
- On QFX3500, QFX3600, QFX5100, and EX4600 switches, if you disable an IRB interface, reboot the switch, and then reenabling the IRB interface, the IRB interface might not be reachable. [PR1196380](#)
- On a Virtual Chassis Fabric, you might see an error such as MMU ERR Type: 1B error, Addr: 0x001052cf, module: 42, which indicates that there was an ECC error in the PFE MMU counter memory. ECC errors are corrected by the hardware without software intervention and are corrected only when a packet hits that memory. Reading an ECC-errored entry always generates an interrupt; however, the error will only be corrected when the packet hits the memory. Because this is a counter memory, the

counter thread reads this memory continuously, and hence you see continuous error messages. [PR11968162](#)

- On QFX5100 and QFX10002 switches, **Rx power low warning set** messages might be logged continuously for channelization ports that are in the DOWN state with snmpwalk running in the background. [PR1204988](#)
- There are basically three arguments—periodic, diagnostic, and tx—for the **lcdd_cmd -f 0 -d chassism -c** command, and this top-level command requires different numbers of arguments. If any one of the arguments is missing when the command is executed on a QFX3500 or QFX3600 switch, chassisd might crash. [PR1206328](#)
- On QFX5100 and EX4600 switches, in rare cases, the fxpc process might crash and restart with a core file generated upon LPM route install failure. After the switch restarts, services are restored. [PR1212685](#)

Routing Protocols

- On QFX5100 switches, the routing protocol process (rpd) fails to respond to any new CLI routing commands (for example, **show mpls lsp terse**). The rpd is forking a child process while processing a **show** command. When the subprocess tries to exit, it attempts to close the management socket being used by the **show** command. This failure might cause the rpd subprocess to crash and generate a core file. It also removes the rpd pid file, which prevents the rpd from processing any new CLI commands even though the original rpd process continues to run normally. [PR1111526](#)

Spanning-Tree Protocols

- On QFX5100 and EX4600 switches, in a scenario where MSTP, RSTP, or VSTP is configured to prevent a Layer 2 network loop, xSTP convergence might fail on an interface that is configured with **flexible-vlan-tagging** and encapsulation of **extended-vlan-bridge**. [PR1179167](#)

Virtual Chassis

- On a non-mixed QFX5100 Virtual Chassis Fabric (VCF) or Virtual Chassis, LACP might flap when the switch in the master Routing Engine role is rebooted using the CLI or because of a power cycle. This issue is not experienced after a Routing Engine switchover. As a workaround, configure a slow LACP timeout. [PR1034377](#)
- On a VCF platform, the memory usage limitation for the vccpd process is 131 MB in memory. Any VCP port flapping will cause a small memory leak (256 KB~1 MB) in the VCF. If the memory usage reached is 131 MB, then the vccpd will crash and create a core file and then restart. In the meantime, a member of the VCF will disconnect from VCF; this will have a service impact until the vccpd comes up again. [PR1158798](#)

Resolved Issues: Release 15.1R4

- [Class of Service \(CoS\)](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [Platform and Infrastructure](#)

- [Routing Policy and Firewall Filters](#)
- [Routing Protocols](#)
- [Security](#)
- [Software-Defined Networks \(SDN\)](#)
- [Virtual Chassis and Virtual Chassis Fabric \(VCF\)](#)

Class of Service (CoS)

- On QFX5100 and EX4600 switches, ICMP, SSH, and ARP traffic generated by the switch might be forwarded to queue 7 (network-control); the default behavior is that the traffic would be forwarded to queue 0 (best-effort). [PR1178188](#)

Interfaces and Chassis

- On a QFX5100 Virtual Chassis, if you configure an aggregated Ethernet interface as an OVSDB interface with multiple subinterfaces that are configured under different VXLAN domains, removal of the last but one AE subinterface might reset VXLAN settings on the physical port that are part of the AE interface, resulting in packet drops. [PR1150467](#)
- On QFX Series and EX Series switches, if you configure VRRP with an MC-LAG between the master and backup switches, both VRRP members of IRB interfaces might stay in the master state after a software upgrade. [PR1157075](#)
- On QFX5100 switches, if a trunk interface is a VXLAN port, tagged frames matching the native VLAN ID might be sent out with the native VLAN tagged. [PR1164850](#)
- If a QFX5100 Virtual Chassis is created with a QFX5100-48S in the routing-engine role and a QFX5100-48T in the linecard role, ports of the QFX5100-48T might be shown as having media type Fiber. [PR1166810](#)
- On QFX5100 switches, if you enable aggregated Ethernet links by deleting the **disable** command, LACP core files might be generated. [PR1173562](#)

Layer 2 Features

- On a QFX5100 switch, if you delete a VLAN and create a new VLAN with a different VLAN ID but use the same VNI, and you commit those changes within a single commit, a MAC learning failure might occur on the newly created VLAN. These system logging messages might be displayed:
 - `fpc0 BCM-VIRTUAL,brcm_vxlan_hw_add(),263:Failed to Program vxlan bd(22) token(0xf) status(-8)`
 - `fpc0 BCM-VIRTUAL,brcm_virtual_bd_add(),626:Cannot create Virtual-BD for bd(22)`
 - `fpc0 BCM-VIRTUAL,brcm_virtual_port_add(),101:Port(ge-0/1/2) add came before bd(22) add`
 - `fpc0 LBCM-L2,pfe_bcm_l2_addr_delete_by_vlan(),52:delete L2 entries associated with bd 21(65535) failed(-4)`

[PR1161574](#)

- On QFX5100 and EX4600 switches, every time a MAC address is learned, some messages might be output to syslog and be repeated frequently. The logged messages have no impact on service traffic. [PR1171523](#)

Platform and Infrastructure

- On QFX Series mixed Virtual Chassis Fabric (VCF), software rollback with the force option (**request system software rollback force**) might not work. [PR1028666](#)
- In a Virtual Chassis Fabric (VCF) with three or four spine devices, the spine devices operating in the linecard role cannot assume the Routing Engine role, including in cases where the master or backup Routing Engine fails. [PR1115323](#)
- In a Virtual Chassis or a Virtual Chassis Fabric (VCF), issuing the **clear arp** command might not clear ARP entries. [PR1159447](#)
- If DHCP packets with MPLS tags are sent to the CPU on a QFX5100 node acting as a PHP node, the logical interfaces index on the packet notification might not be set correctly, and the DHCP packets might be dropped. [PR1164675](#)
- On a QFX5100 switch with an integrated routing and bridging (IRB) interface configured as a Layer 3 interface and with two hosts (Host A and Host B) connected to the switch, if you deactivate the IP address on Host A and then configure the same IP address on Host B, the outgoing interface of the IP address might not be changed in the ARP table. [PR1166400](#)
- Some interfaces might be down after you disable and then reenables autonegotiation on QFX5100 48T-6Q interfaces that are connected to QFX3500 SFP-T interfaces. As a workaround, restart the Packet Forwarding Engine. [PR1168581](#)

Routing Policy and Firewall Filters

- On QFX5100 switches, starting with Junos OS Release 15.1R1, **forwarding-class mcast** configurations are not supported in port-based firewall filters. [PR1088313](#)

Routing Protocols

- On QFX Series switches, when a neighbor device sends a flood of Link Layer Discovery Protocol (LLDP) traffic bigger than 1000 pps to the QFX, Link Aggregation Control Protocol (LACP) flaps might be seen on unrelated interfaces. [PR1058565](#)
- On QFX5100 and EX4600 switches, if you use the Network Configuration Protocol (NETCONF) to add or delete firewall filters on an integrated bridging and routing (IRB) interface, the Packet Forwarding Engine Manager (fxpc) might generate a core file. [PR1155692](#)
- On QFX5100 and EX4600 switches, when a limit traffic filter is configured with TTL=1 packets accepted on the loopback interface, the host-bound unicast packets with TTL=1 (for example, OSPF packets) might be dropped. [PR1161936](#)

- On a QFX3500 switch, if you configure one interface with PIM and the interface sends hello packets, and then you change its PIM hello-interval from non-zero to 0, the interface sends hello packets continuously. [PR1166236](#)
- On QFX5100 switches, if you apply a firewall filter on the loopback interface with the match condition for packets with TTL 0/1 and with **policer** set as the action, the term does not catch the packets. [PR1166936](#)

Security

- On QFX Series switches, up to four port-mirroring analyzers can be configured, which can have up to four ingresses and egresses total for all input stanzas. If the count of ingresses plus egresses is greater than four, the analyzers do not work properly. [PR1168528](#)

Software-Defined Networks (SDN)

- On QFX5100 switches, the openflowd process might generate a core file. [PR1142563](#)

Virtual Chassis and Virtual Chassis Fabric (VCF)

- On QFX5100 Virtual Chassis, if you insert some SFP or SFP+ optics in a port, that port might go down and might not read any other optics. As a workaround, reboot the chassis. [PR1144190](#)
- On QFX5100 Virtual Chassis, Virtual Chassis ports (VCPs) might not be auto-configured if the ports are connected while other ports are being converted. [PR1159242](#)
- On an EX4600 Virtual Chassis or a QFX Series Virtual Chassis or Virtual Chassis Fabric (VCF), if you convert the Virtual Chassis port (VCP) to a network port by issuing the **request virtual-chassis vc-port delete** command, broadcast and multicast traffic might be dropped due to the port remaining programmed as a VCP in the hardware. [PR1159461](#)

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 this section.

- [Authentication and Access Control](#)
- [Bridging and Learning](#)
- [Class of Service \(CoS\)](#)
- [Dynamic Host Control Protocol](#)
- [Firewall Filters](#)
- [High Availability \(HA\) and Resiliency](#)
- [Infrastructure](#)
- [Interfaces and Chassis](#)
- [Layer 2 Features](#)
- [MPLS](#)

- [Multicast](#)
- [Platform and Infrastructure](#)
- [Routing Protocols](#)
- [Software-Defined Networks \(SDN\)](#)
- [Spanning-Tree Protocols](#)
- [Storage and Fibre Channel](#)
- [Virtual Chassis and Virtual Chassis Fabric \(VCF\)](#)

Authentication and Access Control

- On EX4300, EX4600, EX9200, and QFX5100 switches configured for 802.1X authentication, if the VLAN assigned to an access port is changed, then the supplicants authenticated are disconnected and the users are not able to authenticate anymore. [PR1148486](#)

Bridging and Learning

- On EX4300 and QFX Series switches with PVLAN configured, if secondary VLANs (isolated VLANs or community VLANs) are configured with vlan-name, after binding or unbinding the isolated or community VLANs in the primary VLAN, packet loss might occur between existing VLANs. [PR1144667](#)

Class of Service (CoS)

- On QFX Series switches with Data Center Bridging and Capability Exchange (DCBX) enabled, when you are configuring a guaranteed minimum rate of transmission for a CoS traffic control profile, the Layer 2 Control Protocol daemon (l2cpd) might crash during the initial LACP setup. [PR1143216](#)
- On EX4600 and QFX5100 switches, when the Virtual Router Redundancy Protocol (VRRP) priority is modified to change the VRRP mastership after cosd restart (or device restart), packets might be dropped on interfaces that have both inet and inet6 families enabled. [PR1105963](#)
- On QFX5100 and EX4600 switches, if you channelize a 40-Gigabit Ethernet QSFP+ interface into four 10-Gigabit Ethernet ports and try to apply the CoS configuration to one of the specific channels, multicast traffic might get dropped. [PR1108103](#)
- On QFX5100 and EX4600 switches, if an interface that is enabled for flow control is connected to an EX Series switch (except EX9200), even low-rate traffic (host-bound traffic) received might cause a MAC pause frame to be sent from the interface to the peer device, and other transmitting traffic from the interfaces might be affected (for example, LACP flapping might occur). [PR1113937](#)

Dynamic Host Control Protocol

- On QFX5100 switches that are configured with the **include-option-82 nak** option so that Dynamic Host Configuration Protocol (DHCP) servers include option 82 information in NAK messages, two copies of option-82 might be appended to DHCP ACK packets. [PR1064969](#)

- On EX9200 and QFX5100 switches, when DHCP relay is configured with the DHCP server and DHCP client in separate routing instances, unicast DHCP reply packets, for example, DHCPACK in response to a lease renewal request, might be dropped. [PR1079980](#)
- On an EX Series or QFX Series switch configured as a DHCP client, the length of the DHCP vendor ID is always 60 in DHCP discover packets when the vendor class ID is configured, although the actual vendor-id name is less than 60. As per RFC 2132, the code for this option ("Vendor class identifier") is 60, and its minimum length is 1. [PR1123111](#)

Firewall Filters

- On EX4600 and QFX Series switches, if filter-based forwarding (FBF) is configured on an IRB interface that is also enabled for Virtual Router Redundancy Protocol (VRRP), when the host uses the VIP address as the gateway, the switch does not forward packets from that host to the destination routing instance through FBF. This is expected behavior based on the implementation of family inet filters. As a workaround, configure the hosts to use the physical IP address of the IRB interface rather than the VRRP VIP address as the gateway. [PR1025312](#)
- On QFX5100 switches with DHCP relay enabled, if there is a firewall filter with the term "then log" configured, DHCP clients might fail to get IP addresses from a DHCP server. This occurs because the DHCP-relay traffic on the switch drops as the result of rate-limiting. [PR1041513](#)
- On EX4600 and QFX Series switches, you might not be able to commit the configuration when the arp-type match condition is configured in a firewall filter. [PR1084579](#)
- On QFX5100 switches, in the absence of any match condition in filters used for filter-based forwarding (FBF) that are applied to IPv4 traffic, IPv6 traffic coming in on the same interface might get filtered as well. [PR1145667](#)

High Availability (HA) and Resiliency

- On QFX5100 switches with a minimum interval for a Bidirectional Forwarding Detection (BFD) session configured to less than a second, the pre-ISSU check might be successful and continue to implement the ISSU, causing the BFD session to flap. The expected behavior is that the pre-ISSU check for the BFD session fails and ISSU is aborted. [PR1132797](#)

Infrastructure

- On QFX3500, QFX3600, and QFX5100 switches, when **family ethernet-switching** is configured on an interface that is also configured with **encapsulation extended-vlan-bridge**, then transit packets (for example, IP, ping, or Q-in-Q packets) might be dropped on this interface. [PR1078076](#)
- On a QFX3500 switch with nonstop active routing (NSR) enabled, deleting a routing-instance or logical-system configuration might cause a soft assert of the rpd process. If NSR is not enabled, after you delete a routing-instance or logical-system configuration, executing the **restart routing** command might trigger this issue, too. This issue has no functional impact. [PR1102767](#)

- On a Virtual Chassis formed with QFX3500 and QFX3600 switches, CPU consumption might be high if a greater than usual amount of host traffic goes to a VRRP backup node. [PR1124038](#)

Interfaces and Chassis

- On QFX5100 switches, the maximum number of LAGs is now 1000. [PR1082043](#)
- On a QFX5100 Virtual Chassis, the MAC address is not learned on an aggregated Ethernet (AE) interface configured as a VXLAN Layer 2 port and with the interface mode configured as access. The issue is observed only with AE interfaces that span multiple Virtual Chassis members and when the member node is rebooted or power cycled. [PR1112790](#)
- On QFX5100 switches, if an mc-ae member link is deleted and then re-added on an MC-LAG node, there could be a traffic loss of about 2 seconds. [PR1146206](#)
- On QFX5100 switches, a child member might drop the incoming Link Aggregation Control Protocol (LACP) frames when this child member is moved from an access-mode VXLAN LAG interface to a trunk-mode VXLAN LAG interface. [PR1153042](#)
- On QFX5100 and EX4600 switches, the Gigabit Ethernet (ge) interface might stop forwarding traffic when you hot-swap a transceiver from SFP-SX to SFP-T. [PR1144485](#)

Layer 2 Features

- On QFX5100 and EX4600 switches running under Junos OS Release 14.1X53-D10 or later, when DHCPv6 solicitation packets go through the device with Q-in-Q configured, the packets might be dropped by peers because the S-tag has not been added. [PR1103793](#)
- On EX4300, EX4600, and QFX Series switches, if a trunk port is deleted and then reconfigured as an access port in the same commit, the Layer 2 address learning daemon (l2ald) might generate a core file. [PR1105255](#)
- On EX4600 and QFX5100 switches, the VLAN Spanning Tree Protocol (VSTP) bridge protocol data units (BPDUs) might be reinjected to the Packet Forwarding Engine and not be sent out of an interface when the interface has been added to the VSTP configuration and is configured with flexible-vlan-tagging. [PR1117540](#)
- On QFX5100 switches, if you configure a PVLAN inter-switch link on an existing working trunk port, normal VLAN traffic might break. [PR1118728](#)
- On EX4300, EX4600, and QFX Series switches, traffic received on the backup redundant trunk group (RTG) link might get forwarded to other interfaces following an RTG link failover. [PR1119654](#)
- If you reboot one FPC in a two-member Virtual Chassis, the traffic might not exit from the FPC after the FPC comes back online and rejoins the Virtual Chassis, and local registers might be incorrectly cleared if the port number is the same on both the master and backup. [PR1124162](#)

- On a QFX5100 Virtual Chassis, traffic might not pass the inter-member when the firewall filter is applied to the ingress interface using the interface **vlan** option. [PR1138714](#)
- On QFX5100 and EX4600 switches, after you delete one logical interface from one VLAN that is configured with multiple logical interfaces, the MAC address for other logical interfaces might not be learned again. [PR1149396](#)

MPLS

- On QFX5100 switches, a ping from the customer edge (CE) to the provider edge (PE) (last-hop router [LHR]) lo0 interface does not go through with explicit-null (RSVP). [PR1145437](#)

Multicast

- On EX4600 and QFX Series switches, IGMP snooping might not be enabled after you reboot the switch. You might see the same issue after you run a nonstop software upgrade (NSSU) on the switch. [PR1082453](#)

Platform and Infrastructure

- Setting link speed to 100 Mbps does not work in the following situations:
 - When network interfaces are used on an EX4600
 - When an EX4600-EM-8F expansion module is installed in a QFX5100-24Q switch or EX4600 switch

[PR1032557](#)

- On EX Series and QFX Series switches, issuing the **show interfaces extensive** command or polling SNMP OID ifOutDiscards provides a drop count of zero. [PR1071379](#)
- On QFX5100 switches, the wrong source IP address is being used when the switch initiates traffic and em0 is configured with a 192.168.1.x/16 subnet and after the switch has been upgraded with the force-host option. [PR1071517](#)
- On EX4600 and QFX Series switches, MAC addresses on one VLAN might be installed in the hardware but be missing from the Ethernet-switching table if the following steps were taken and if $A + B \geq 4096$:
 1. Configured **vlan-id-list** for a VLAN range "A" with a commit.
 2. Deleted the VLAN range "A" and re-added the VLAN range "B" in the same commit.

[PR1074919](#)

- On QFX3500 switches, if you remove 1-Gigabit Ethernet SFP transceivers from ports 0-5/42-47 and then insert 10-Gigabit Ethernet SFP+ transceivers in the same ports, the 10GE SFP+ transceivers might not be detected. [PR1085634](#)
- On QFX5100 switches, adding or removing virtual routing and forwarding (VRF) instances that have many logical interfaces in the link aggregation group (LAG) might cause Link Aggregation Control Protocol (LACP) flapping. [PR1087615](#)
- On EX4600 and QFX5100 switches, when Spanning Tree Protocol (STP) is enabled on an S-VLAN, that S-VLAN's STP bridge data protocol unit (BPDU) packets might

be dropped by the S-VLAN interface if the S-VLAN interface is an aggregated Ethernet (AE) interface. [PR1089331](#)

- On a QFX Series Virtual Chassis Fabric (VCF), rebooting a leaf node might change the size of the VCF, resulting in a flood loop of the unicast or multicast traffic. To fix the issue, use the new configuration statement `fabric-tree-root`. [PR1093988](#)
- On EX4600 and QFX5100 switches, when flow control is configured on an interface, and pause frames are sent to this interface, the interface might go down. [PR1098055](#)
- On QFX Series switches, removing or inserting one QSFP might cause the pfe process to crash. [PR1098385](#)
- On EX4600 and QFX5100 switches with Q-in-Q, if the native VLAN is configured on a Q-in-Q interface connected to a customer edge (CE), the packets going out with the native VLAN ID (customer-VLAN) are still tagged. [PR1105247](#)
- On a QFX Series Virtual Chassis Fabric (VCF) or Virtual Chassis with graceful Routing Engine switchover (GRES) enabled, the backup Routing Engine might continuously reboot after you configure `forward-and-send-to-re` or `forward-only` under the `[edit interface interface-name unit unit-number family inet targeted-broadcast]` hierarchy. [PR1106151](#)
- On a QFX5100 VCF in auto-provisioned mode, when adding a new leaf device to the VCF, you should zeroize the device and reboot by using the `request system zeroize` command if the new leaf device has been configured with any command. The issue (interface still up) might be observed at the time of the reboot until the Packet Forwarding Engine reinitializes the interfaces. [PR1106194](#)
- On EX4300 and QFX Series switches, the analytics daemon (analyticsd) runs on devices even if there is no analytics configuration, which might cause system instability because of the high number of files opened by analyticsd. [PR111613](#)
- On QFX5100 Virtual Chassis, multiple PFEMAN disconnects and reconnects between the master and backup within a short period of time can cause the backup to generate core files. [PR1123379](#)
- On EX4300, EX4600, EX9200, and QFX Series switches, the lldp-med-bypass feature does not work. [PR1124537](#)
- On QFX3500 and QFX5100 switches, if you commit an `et inet` interface with an MPLS configuration and the `no-redirects` statement, the operation might cause no protocol ARP for the specific logical interface in the Packet Forwarding Engine, and traffic is not sent out. [PR1138310](#)
- On QFX Series and EX4600 switches, if an aggregated Ethernet (AE) interface is used as an ECMP next hop (load balance), traffic is not hashed evenly to all member interfaces correctly. [PR1141571](#)
- On EX4200, EX4300, EX4550, EX4600, and QFX5100 switches with Media Access Control Security (MACsec) enabled on an AE subinterface, MACsec might not work because the MACsec Key Agreement (MKA) session is not established with a peer after `flexible-vlan-tagging` is configured on the AE interface. [PR1133528](#)

- On QFX5100 switches, if you delete an autonegotiate configuration on a 10-gigabit interface (xe), the interface goes down as expected because the autonegotiate setting is not matching with that on the peer interface. However, the interface might come up after the reboot even though autonegotiate is still disabled. [PR1144718](#)
- On EX Series and QFX Series switches, if **interface-mac-limit** is configured on an interface range, the commit might fail. [PR1154699](#)

Routing Protocols

- On a standalone QFX Series switch, if you configure a nested firewall filter and then attempt to commit the configuration, the firewall compiler process (dfwc) might crash and generate a core file, leading to commit failure. [PR1094428](#)
- On a QFX VCF, if the switch works as part of a target subnet, while receiving the targeted broadcast traffic, the packets might be forwarded to the destination with the switch's MAC address as the destination MAC address, where it should be converted into a Layer 2 broadcast frame with destination MAC address FFFF.FFFF.FFFF. [PR1114717](#)
- On QFX5100 switches, you might see the **soc_mem_read: invalid index -1 for memory EGR_L3_INTF** log message. You can ignore the message; there is no functional impact on the switch. [PR1126035](#)

Software-Defined Networks (SDN)

- In an OpenFlow scenario with QFX5100 or EX9200 as the virtual switch, the openflowd process might crash after you issue the **show openflow statistics tables** command. [PR1131697](#)

Spanning-Tree Protocols

- On QFX5100 switches, when an STP configuration is initially applied to an interface and the interface is down at that moment, executing **show** or **clear spanning-tree statistic interface** might cause the Layer 2 control protocol process (l2cpd) to crash. [PR1152396](#)

Storage and Fibre Channel

- On EX4500 and QFX Series switches with Data Center Bridging Capability Exchange (DCBX) enabled, when the DCBX neighbor is up and then receives a normal Link Layer Discovery Protocol (LLDP) packet (without DCBX TLVs) on the same port as the DCBX packets, the device might ignore the DCBX packets, causing session timeouts and a reset of the priority-based flow control (PFC) settings. [PR1095265](#)

Virtual Chassis and Virtual Chassis Fabric (VCF)

- On a Virtual Chassis Fabric (VCF), a small amount of Layer 3 unicast packet loss (for example, 0.2 - 0.3 sec) might be seen when a leaf node that is not in the traffic path is rebooted. [PR976080](#)
- On a QFX Series Virtual Chassis Fabric (VCF), rebooting a leaf node might change the size of the VCF, resulting in a flood loop of the unicast or multicast traffic. To fix the issue, use the new configuration statement **fabric-tree-root**. [PR1093988](#)

- See Also**
- [New and Changed Features on page 366](#)
 - [Changes in Behavior and Syntax on page 369](#)
 - [Known Behavior on page 370](#)
 - [Known Issues on page 374](#)
 - [Documentation Updates on page 390](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 390](#)
 - [Product Compatibility on page 394](#)

Documentation Updates

There are no errata or changes in Junos OS Release 15.1R6 for the QFX Series switches documentation.

- See Also**
- [New and Changed Features on page 366](#)
 - [Changes in Behavior and Syntax on page 369](#)
 - [Known Behavior on page 370](#)
 - [Known Issues on page 374](#)
 - [Resolved Issues on page 375](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 390](#)
 - [Product Compatibility on page 394](#)

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

- [Upgrading Software on QFX3500, QFX3600, and QFX5100 Standalone Switches on page 390](#)
- [Performing an In-Service Software Upgrade \(ISSU\) on the QFX5100 Switch on page 392](#)

Upgrading Software on QFX3500, QFX3600, and QFX5100 Standalone Switches

When upgrading or downgrading Junos OS, always use the jinstall package. Use other packages (such as the jbundle package) only when so instructed by a Juniper Networks support representative. For information about the contents of the jinstall package and details of the installation process, see the [Junos OS Installation and Upgrade Guide](#) and [Junos OS Basics](#) in the QFX Series documentation.

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to <https://support.juniper.net/support/downloads/junos.html> .
The Junos Platforms Download Software page appears.
2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.
3. Select **15.1** in the Release pull-down list to the right of the Software tab on the Download Software page.
4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 15.1 release.
An Alert box appears.
5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.
A login screen appears.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Download the software to a local host.
8. Copy the software to the device or to your internal software distribution site.
9. Install the new jinstall package on the device.



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

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

```
user@host> request system software add source/jinstall-qfx-5-15.1-R3-domestic-signed.tgz
reboot
```

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

- **/pathname**—For a software package that is installed from a local directory on the switch.
- For software packages that are downloaded and installed from a remote location:
 - **ftp://hostname/pathname**
 - **http://hostname/pathname**

- `scp://hostname/pathname` (available only for Canada and U.S. version)

Adding the **reboot** command reboots the switch after the upgrade is installed. When the reboot is complete, the switch displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.



NOTE: After you install a Junos OS Release 15.1 jinstall package, you can issue the **request system software rollback** command to return to the previously installed software.

Performing an In-Service Software Upgrade (ISSU) on the QFX5100 Switch

You can use ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.



NOTE: ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Before you begin software installation using ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

To verify that nonstop active routing is enabled:



NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

```
user@switch> show task replication
Stateful Replication: Enabled
RE mode: Master
```

If nonstop active routing is not enabled (**Stateful Replication is Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the **request system snapshot** command.

To upgrade the switch using ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the `/var/tmp` directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
 - On the switch, enter:

```
user@switch> request system software in-service-upgrade
/var/tmp/package-name.tgz
```

where *package-name.tgz* is, for example, *jinstall-132_x51_vjunos.domestic.tgz*.



NOTE: During the upgrade, you cannot access the Junos OS CLI.

The switch displays status messages similar to the following messages as the upgrade executes:

```
warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
lost!
ISSU: Validating Image
ISSU: Preparing Backup RE
Prepare for ISSU
ISSU: Backup RE Prepare Done
Extracting jinstall-qfx-5-13.2X51-D15.4-domestic ...
Install jinstall-qfx-5-13.2X51-D15.4-domestic completed
Spawning the backup RE
Spawn backup RE, index 0 successful
GRES in progress
GRES done in 0 seconds
Waiting for backup RE switchover ready
GRES operational
Copying home directories
Copying home directories successful
Initiating Chassis In-Service-Upgrade
Chassis ISSU Started
ISSU: Preparing Daemons
ISSU: Daemons Ready for ISSU
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item          Status          Reason
  FPC 0         Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
```

ISSU: IDLE
Initiate em0 device handoff



NOTE: An ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).



NOTE: If the ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at `/var/log/vjunos-log.tgz`.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

```
user@switch> show version
```

6. To ensure that the resilient dual-root partitions feature operates correctly, copy the new Junos OS image into the alternate root partitions of all of the switches:

```
user@switch> request system snapshot slice alternate
```

Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

- See Also**
- [New and Changed Features on page 366](#)
 - [Changes in Behavior and Syntax on page 369](#)
 - [Known Behavior on page 370](#)
 - [Known Issues on page 374](#)
 - [Resolved Issues on page 375](#)
 - [Documentation Updates on page 390](#)
 - [Product Compatibility on page 394](#)

Product Compatibility

- [Hardware Compatibility on page 394](#)

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 QFX 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:
<https://pathfinder.juniper.net/feature-explorer/>

- See Also**
- [New and Changed Features on page 366](#)
 - [Changes in Behavior and Syntax on page 369](#)
 - [Known Behavior on page 370](#)
 - [Known Issues on page 374](#)
 - [Resolved Issues on page 375](#)
 - [Documentation Updates on page 390](#)
 - [Migration, Upgrade, and Downgrade Instructions on page 390](#)

Third-Party Components

This product includes third-party components. To obtain a complete list of third-party components, see [Overview for Routing Devices](#).

For a list of open source attributes for this Junos OS release, see [Open Source: Source Files and Attributions](#).

Upgrading Using Unified ISSU

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified in-service software upgrade is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified in-service software upgrade, see the [High Availability Feature Guide for Routing Devices](#).

For information about unified ISSU support across platforms and Junos OS releases, see the [In-Service Software Upgrade \(ISSU\)](#) Web application.

Compliance Advisor

For regulatory compliance information about [Common Criteria](#), [FIPS](#), [Homologation](#), [RoHS2](#), and [USGv6](#) for Juniper Networks products, see the [Juniper Networks Compliance Advisor](#).

Finding More Information


For the latest, most complete information about known and resolved issues with Junos OS, see the Juniper Networks Problem Report Search application at: <https://prsearch.juniper.net>.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at: <https://pathfinder.juniper.net/feature-explorer/>.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the [Juniper Networks TechLibrary](#) site, and do one of the following:



- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.
- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at <https://www.juniper.net/assets/us/en/local/pdf/resource-guides/7100059-en.pdf>.
- Product warranties—For product warranty information, visit <https://support.juniper.net/support/warranty/>.
- JTAC hours of operation—The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: <https://www.juniper.net/customers/support/>
- Search for known bugs: <https://www2.juniper.net/kb/>
- Find product documentation: <https://www.juniper.net/documentation/>
- Find solutions and answer questions using our Knowledge Base: <https://kb.juniper.net/>
- Download the latest versions of software and review release notes: <https://www.juniper.net/customers/csc/software/>
- Search technical bulletins for relevant hardware and software notifications: <https://kb.juniper.net/InfoCenter/>
- Join and participate in the Juniper Networks Community Forum: <https://www.juniper.net/company/communities/>
- Open a case online in the CSC Case Management tool: <https://www.juniper.net/cm/>

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: <https://entitlementsearch.juniper.net/entitlementsearch/>

Opening a Case with JTAC

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

- Use the Case Management tool in the CSC at <https://www.juniper.net/cm/>.

- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

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

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

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

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

Revision History

4 October 2018—Revision 17, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

25 January 2018—Revision 16, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

7 December 2017—Revision 15, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

26 October 2017—Revision 14, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

30 June 2017—Revision 13, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

29 June 2017—Revision 12, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

23 June 2017—Revision 11, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

22 June 2017—Revision 10, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

15 June 2017—Revision 9, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

9 June 2017—Revision 8, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

8 June 2017—Revision 7, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

1 June 2017—Revision 6, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

25 May 2017—Revision 5, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

19 May 2017—Revision 4, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

11 May 2017—Revision 3, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

4 May 2017—Revision 2, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

27 April 2017—Revision 1, Junos OS Release 15.1R6— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

23 February 2017—Revision 6, Junos OS Release 15.1R5— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

19 January 2017—Revision 5, Junos OS Release 15.1R5— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

12 January 2017—Revision 4, Junos OS Release 15.1R5— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

14 December 2016—Revision 3, Junos OS Release 15.1R5— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

7 December 2016—Revision 2, Junos OS Release 15.1R5— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

30 November 2016—Revision 1, Junos OS Release 15.1R5— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

27 October 2016—Revision 7, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

17 August 2016—Revision 6, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

4 August 2016—Revision 5, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

15 July 2016—Revision 4, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

14 July 2016—Revision 3, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

7 July 2016—Revision 2, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

30 June 2016—Revision 1, Junos OS Release 15.1R4— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

10 June 2016—Revision 8, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

26 May 2016—Revision 7, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

19 May 2016—Revision 6, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

12 May 2016—Revision 5, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

21 April 2016—Revision 4, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

13 April 2016—Revision 3, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

6 April 2016—Revision 2, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

30 March 2016—Revision 1, Junos OS Release 15.1R3— ACX Series, EX Series, M Series, MX Series, PTX Series, QFX Series, and T Series.

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

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

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

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

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

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

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

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

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

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

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

5 June 2015—Revision 1, Junos OS Release 15.1R1— EX Series, M Series, MX Series, PTX Series, and T Series.

Copyright © 2018 Juniper Networks, Inc. All rights reserved.

Juniper Networks, the Juniper Networks logo, Juniper, and Junos are registered trademarks of Juniper Networks, Inc. and/or its affiliates in the United States and other countries. All other trademarks may be property of their respective owners.

Juniper Networks assumes no responsibility for any inaccuracies in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.