Release Notes: Junos® OS Release 16.2R2 for the ACX Series, EX Series, MX Series, PTX Series, and T Series

18 July 2019

Contents

Introduction

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

These release notes accompany Junos OS Release 16.2R2 for the ACX Series, EX Series, MX Series, PTX Series, and T Series. . They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

These release notes accompany Junos OS Release 16.2R2 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/software/junos/.

New and Changed Features

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for ACX Series Universal Metro Routers.

Release 16.2R2 New and Changed Features

This section describes the new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 16.2R2.

Class of Service

- **Support for CoS (ACX5000)**—Junos OS for ACX5000 line of routers supports the following class of service (CoS) features:
  - Ingress classification
  - Fixed classification (interface-based classification)

- Behavior aggregate (BA) classification
- Multifield (MF) classification

CoS features include rewrite, scheduling and buffer management, host outbound traffic, and statistics.

In addition to these features, you can configure buffer partitions for multicast packets and shared buffer that the multicast packets in the queue can consume. To configure these features, use the **buffer-partition multicast percent** and **multicast** statements at the **[edit class-of-service schedulers]** hierarchy level.

The following CoS behaviors are specific to the ACX5000 line of routers:

- Strict priority queuing—Unlike other ACX routers, ACX5000 line of routers support committed information rate (CIR) among strict-priority queues. There is no implicit queue number-based priority among the strict-priority queues.
- Weighted random early detection (WRED)—Unlike other ACX routers, ACX5000 line of routers support configuring drop profiles (to specify different drop behavior) for loss priorities low, medium-high, and high for both TCP and non-TCP protocols.

- **Support for IPv6 CoS (ACX5000)**—Junos OS for ACX5000 line of routers supports IPv6 (**dscp-ipv6**) classification and rewrite.

Firewall

- **Support for IPv6 firewall filter (ACX5000)**—Junos OS for ACX5000 line of routers supports IPv6 firewall filter at the [**edit firewall family inet6 filter *filter-name***] hierarchy level.

- **Support for CoS, filter, and policer with VPLS (ACX5000)**—Junos OS for ACX5000 line of routers supports Class of Service (CoS), firewall filters, and policers with the VPLS feature. The ACX5000 line of routers support CoS ingress classification and egress rewrite features with VPLS. VPLS firewall filters and policers can be configured at the logical interface family level.

IPv6

- **Support for IPv6 VPN provider edge router (6VPE) over MPLS (ACX5000)**—Junos OS for ACX5000 line of routers provides IPv6 VPN provider edge router (6VPE) support over MPLS. ACX5000 line of routers act as a VPN provider edge router that provides IPV6 forwarding over MPLS. 6VPE adds IPv6 support to the current IPv4 MPLS by transporting IPv6 across MPLS core.

- **Support for DHCPv6 relay agent (ACX5000)**—Junos OS for ACX5000 line of routers supports DHCPv6 relay agent. The DHCPv6 relay agent enhances the extended DHCP relay agent by providing DHCP support in an IPv6 network. DHCPv6 relay agents eliminate the necessity of having a DHCPv6 server on each physical network. An ACX5000 router configured as a DHCPv6 relay agent passes messages between the DHCPv6 client and the DHCPv6 server, similar to the way a DHCP relay agent supports an IPv4 network.

To configure the DHCPv6 relay agent on ACX5000 line of routers, include the **dhcpv6** statement at the [**edit forwarding-options dhcp-relay**] hierarchy level. You can also include the **dhcpv6** statement at the [**edit routing-instances routing-instance-name forwarding-options**] hierarchy level.

- **Support for DHCPv6 server (ACX5000)**—Junos OS supports configuring ACX5000 line of routers as a DHCPv6 server. The DHCPv6 server provides IPv6 address and other configuration information for the clients to configure itself. To configure the DHCPv6 server on the router, include the **dhcpv6** statement at the [**edit system services dhcp-local-server**] hierarchy level.

  To configure the DHCPv6 server in a routing instance, include the **dhcpv6** statement at the [**edit routing-instances routing-instance-name system services dhcp-local-server**] hierarchy level.

  You must also configure individual address pools and the DHCP attributes for the common address pools at the [**edit access**] hierarchy level.

  To maintain DHCPv6 subscribers whenever an interface delete event occurs (such as PFE reboot or crash), the following CLI is supported:

  ```
  [edit system services]
  subscriber-management {
  maintain-subscriber {
  }
  }
  ```

Management

- **Support for Ethernet ring protection switching (ACX5000)**—Junos OS for ACX5000 line of routers supports Ethernet ring protection switching (G.8032v2). With the G.8032v2 capability, the ACX5000 line of routers support manual commands (force switch, manual switch, and clear) and interconnection of multiple Ethernet rings without virtual channels.

- **Support for OAM features (ACX5000)**—Junos OS for ACX5000 line of routers supports the following OAM features:

  - Ethernet OAM
    - IEEE 802.3ah link fault management
    - Connectivity fault management (CFM) of down MEPs and up MEPs
    - ITU Y.1731 delay measurement and synthetic loss measurement (SLM)
  - Virtual circuit connection verification (VCCV) and Bidirectional Forwarding Detection (BFD)

- **Support for Layer 2, IP, MPLS, CoS, firewall, and OAM features (ACX5000)**—Junos OS for ACX5000 line of routers supports Layer 2, IP, MPLS, multicast, CoS, firewall, and OAM features. The ACX5000 line of routers do not support the following features:

- T1/E1 interfaces
- IPsec and NAT services
- Hierarchical policer
- RFC 2544 generator
- Real-time performance monitoring and Two-Way Active Measurement Protocol
- Precision Timing Protocol (PTP) and Synchronized Ethernet

- **Connectivity fault management support for maintenance association intermediate points (ACX5000)**—Junos OS for ACX5000 line of routers supports connectivity fault management (CFM) support for maintenance association intermediate points (MIPs). A MIP provides monitoring capability of intermediate points within a service.

- **Support for analyzer and flow mirroring (ACX5000)**—Junos OS for ACX5000 line of routers supports both port mirroring and analyzer for mirroring the packets received or sent from a specific port. This feature is useful for debugging network problems and to prevent attacks on a network.

  NOTE: **ACX5000 line of routers supports only ingress mirroring. Analyzer with ingress interface or interface list is not supported.**

  The ACX5000 line of routers supports the following mirroring:

  - VLAN mirroring—Support for VLAN mirroring using analyzer where input to mirror is a VLAN (Bridge domain).
  - Flow mirroring—Support for flow-based mirroring where the input for mirror is through firewall filter match and supports only Ethernet-switching and inet family types.

- **Support for unified forwarding table (ACX5000)**—Junos OS for the ACX5000 line of routers supports the use of a unified forwarding table to optimize address storage. Using this feature, you can control the allocation of forwarding table memory available to store the following entries:

  - MAC addresses
  - Layer 3 host entries
  - Longest prefix match (LPM) table entries

  You can use five predefined profiles (**l2-profile-one, l2-profile-two, l2-profile-three, l3-profile, lpm-profile**) to allocate the table memory space differently for each of these entries. You configure and select the profiles that best suits your network environment needs.

  In addition to interface statistics, the following statistics are also supported on the ACX5000 line of routers with increased scale:

  - Logical interface statistics
  - MPLS unicast next hops statistics

- Multicast route statistics

Routing Protocols

- **Support for Virtual Router Redundancy Protocol version 3 (ACX5000)**—Junos OS for ACX5000 line of routers supports Virtual Router Redundancy Protocol (VRRP) version 3. With version 3, VRRP is supported over IPv6 addresses.

  VRRPv3 on ACX5000 line of routers supports:

  - Fast interval with minimum advertisement interval of 100 milliseconds
  - IRB interfaces
  - Aggregated Ethernet (AE) interfaces
  - Auto-generation of link local address

VPLS

- **Support for virtual private LAN service (ACX5000)**—Junos OS for ACX5000 line of routers support the virtual private LAN service (VPLS) feature. With this feature, you can deploy the ACX5000 line of routers as part of a full-mesh VPLS domain, as well as a hub site for hierarchical VPLS (H-VPLS).

  

  NOTE: Applying a forwarding table filter to a VPLS routing instance is not supported on ACX5000 line of routers.

Release 16.2R1 New and Changed Features

This section describes the new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 16.2R1.

Firewall

- **Enhancements to firewall filters (ACX Series)**—Starting in Release 16.2, Junos OS for ACX Series Universal Metro Routers supports the following firewall filter enhancements:
  - Loopback filter support in egress direction.
  - Firewall filter rule match for **source-prefix-list** and **destination-prefix-list**.
  - Additional firewall filter actions on IRB interfaces.
- **Enhancements to support log and syslog firewall filter actions (ACX Series)**—Starting in Release 16.2, Junos OS for ACX Series Universal Metro Routers supports log and syslog firewall filter actions in ingress and egress directions for family **inet** and family **bridge** protocol families.

  The following limitations apply:

- Broadcast, unknown unicast, and multicast (BUM) traffic is not logged for family **bridge** and family **inet** filters in egress direction.
- For egress log and syslog actions, DSCP, TTL, and IEEE 802.1p bits are matched based on ingress values.
- For family**inet**, the log and syslog filter actions do not work at egress if a packet is forwarded through the default route entry to egress.
- For family **bridge**, the log and syslog filter actions do not work at egress if the filter term contains **user-vlan-id**, **user-vlan-pri**, and **user-vlan-dei** match conditions.
- For family **bridge** and family **inet**, if a packet hits log or syslog actions on both the ingress and egress directions, only one log and one syslog message are recorded.
- For family **inet**, if a packet hits reject action on ingress, the packet is not logged on the egress filter action.

- **Enhancements to unicast reverse-path forwarding (uRPF) check (ACX Series)**—Starting in Release 16.2, Junos OS for ACX Series Universal Metro Routers supports uRPF check on IRB interfaces and uRPF fail filter configuration on IPv4 and IPv6 interfaces.

  NOTE: The uRPF fail filter cannot match packets failed at ingress port check (strict mode).

  The uRPF fail filter can match packets failing source IP lookup but cannot match packets failing the input interface check (strict mode).

  The uRPF fail filter applies only to interface-specific instances of the firewall filter.

- **Filter support on the loopback interface (ACX Series)**—Junos OS for ACX Series Universal Metro Routers provide support for applying a firewall filter on the loopback interface (**lo0**). Filters on the loopback interface are applied to protect the Routing Engine from various attacks.

IPv6

- **Support for IPv6 multicast using Multicast Listener Discovery protocol (ACX Series)**—Junos OS for ACX Series Universal Metro Routers support IPv6 multicast using Multicast Listener Discovery (MLD) protocol. To support multicast data delivery, ACX line of routers support MLD (version 1 and version 2) for forming group membership in IPv6 networks and Protocol Independent Multicast (PIM) version 6 to form IPv6 multicast delivery tree.

  To configure MLD, include the **mld** statement at the [**edit protocols**] hierarchy level.

  To configure PIM, include the **pim** statement at the [**edit protocols**] hierarchy level.

See Also
- Changes in Default Behavior and Syntax on page 11

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 16.2R1 for the ACX Series.

General Routing

- **Support for deletion of static routes when the BFD session goes down (ACX Series)**—Starting with Junos OS 16.2R2, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

Interfaces and Chassis

- **Support for logical interfaces**—ACX5048 and ACX5096 routers do not support configuring more than 1000 logical interfaces.

Management

- **Support for status deprecated statement in YANG modules (ACX Series)**—Starting in Junos OS Release 16.1R2, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

Network Management and Monitoring

- **Possible change to the object identifier (ACX Series)**—Starting in Junos OS Release 16.2R2, the many warnings that occurred previously during MIB loading in the SNMP client has been reduced. This update could change the resulting OID depending on the SNMP client-loading logic. For example, jnxProductACX1000 is defined under the two following nodes:
  - jnxProductACX1000 OBJECT IDENTIFIER ::= { jnxProductVariationACX 1 }
  - jnxProductACX1000 OBJECT IDENTIFIER ::= { jnxProductVariationACX1000 1 }

Because the second definition is the duplicate, it is removed. Previously, the SNMP client referred to the second OID based on MIB loading logic and then you would see a change in OID for the client.

Platform and Infrastructure

- **Improvements to MIB validation during Junos OS build (libjsnmp) (ACX Series)**—Many warnings can be ignored if produced while MIBs are compiling. The following warnings should be considered errors because they can break the build:

  [0-9]:.*failed to locate—An OID failed to be located.

  [0-9]:.*redefinition of identifier—Redefinition of OIDs found in jnx-chas-defines.

  [0-9]:.*sequence-type-mismatch—Type mismatch found in sequence syntax of the table and actual OID type.

  [0-9]:.*cannot be imported from module—MIB failed to import because order is not defined properly.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (ACX Series)**—Starting in Junos OS Release 16.2R2, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 16.2R2, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the XML and JSON output when displaying the differences between the candidate and active configurations (ACX Series)**—Starting in Junos OS Release 16.2R2, when you compare the candidate and active configurations and display the differences in XML or JSON format, for example by using the **show | compare | display (json | xml)** CLI command or the **<get-configuration compare="rollback" format="(json | xml)">** RPC, the device omits the **<configuration>** tag in the XML output and omits the **configuration** object in the JSON output if the comparison either returns no differences or if the comparison returns differences for only non-native configuration data, for example, configuration data associated with an OpenConfig data model.

See Also

- New and Changed Features on page 5
- Known Behavior on page 13
- Known Issues on page 13
- Resolved Issues on page 14
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 15
- Product Compatibility on page 15

Known Behavior

There are no known limitations in Junos OS Release 16.2R2 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 5
- Changes in Default Behavior and Syntax on page 11
- Known Issues on page 13
- Resolved Issues on page 14
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 15
- Product Compatibility on page 15

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 16.2R1 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.

- Firewall Filter on page 13

Firewall Filter

- Unexpected firewall filter hits are noticed for a short transient time. This occurs during PFE restart for scaled bridge domain filters having **destination-mac-address** as the match condition. In a dropped packet, if a customer vlan matches with some other service vlan, the log, syslog, or count filter actions will take effect for a short period of time. PR1198151
- ACX hardware supports unicast reverse-path forwarding (uRPF) mode at physical interface level and Junos OS supports uRPF at logical interface level. To avoid the confusion with respect to the uRPF mode, only one mode (strict or loose) should be configured for all the logical interfaces within a physical interface. This also applies to the logical interface in a bridge domain if IRB is configured and uRPF mode is enabled at the logical interface of IRB. PR1196908
- ACX hardware does not support uRPF statistics. The values shown in the Junos OS CLI for uRPF statistics at logical interface level can be ignored. As a workaround, you can use uRPF fail filter configuration where firewal filter has count as the action. The fail filter functionality is limited to loose-mode only. Packets dropped specifically due to strict-mode will not hit the fail-filters. PR1188020

See Also

- New and Changed Features on page 5

- Changes in Default Behavior and Syntax on page 11
- Known Behavior on page 13
- Resolved Issues on page 14
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 15
- Product Compatibility on page 15

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.

See Also

- New and Changed Features on page 5
- Changes in Default Behavior and Syntax on page 11
- Known Behavior on page 13
- Known Issues on page 13
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 15
- Product Compatibility on page 15

Documentation Updates

There are no errata or changes in Junos OS Release 16.2R2 for the ACX Series routers documentation.

See Also

- New and Changed Features on page 5
- Changes in Default Behavior and Syntax on page 11
- Known Behavior on page 13
- Known Issues on page 13
- Resolved Issues on page 14
- Migration, Upgrade, and Downgrade Instructions on page 15
- Product Compatibility on page 15

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 15

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 even from Junos OS Release 14.1 to Release 16.1. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind. For example, you cannot directly upgrade from Junos OS Release 12.1 (a non-EEOL release) to Junos OS Release 13.2 or directly downgrade from Junos OS Release 13.2 to Junos OS Release 12.1.

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

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

See Also

- New and Changed Features on page 5
- Changes in Default Behavior and Syntax on page 11
- Known Behavior on page 13
- Known Issues on page 13
- Resolved Issues on page 14
- Documentation Updates on page 14
- Product Compatibility on page 15

Product Compatibility

- Hardware Compatibility on page 16

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/.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the Hardware Compatibility tool.

See Also

- New and Changed Features on page 5
- Changes in Default Behavior and Syntax on page 11
- Known Behavior on page 13
- Known Issues on page 13
- Resolved Issues on page 14
- Documentation Updates on page 14
- Migration, Upgrade, and Downgrade Instructions on page 15

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 16.2R2 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/software/junos/.

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 19
- Resolved Issues on page 20
- Documentation Updates on page 22
- Migration, Upgrade, and Downgrade Instructions on page 23
- Product Compatibility on page 24

New and Changed Features

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for EX Series.

NOTE: **The following EX Series switches are supported in Release 16.2R2: EX9200.**

Release 16.2R2 New and Changed Features

There are no new features or enhancements to existing features for EX Series in Junos OS Release 16.2R2.

Release 16.2R1 New and Changed Features

There are no new features or enhancements to existing features for EX Series in Junos OS Release 16.2R1.

See Also

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 16.2 for the EX Series.

General Routing

- **Support for deletion of static routes when the BFD session goes down (EX Series)**—Starting in Junos OS Release 16.2R2, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

Management

- **Support for status deprecated statement in YANG modules (EX9200)**—Starting with Junos OS Release 16.2R1, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (EX Series)**—Starting in Junos OS Release 16.2R2, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 16.2R2, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the XML and JSON output when displaying the differences between the candidate and active configurations (EX Series)**—Starting in Junos OS Release 16.2R2, when you compare the candidate and active configurations and display the differences in XML or JSON format, for example by using the **show | compare | display (json | xml)** CLI command or the **<get-configuration compare="rollback" format="(json | xml)">** RPC, the device omits the **<configuration>** tag in the XML output and omits the **configuration** object in the JSON output if the comparison either returns no differences or if the comparison returns differences for only non-native configuration data, for example, configuration data associated with an OpenConfig data model.

See Also

- New and Changed Features on page 17
- Known Behavior on page 18
- Known Issues on page 19
- Resolved Issues on page 20
- Documentation Updates on page 22
- Migration, Upgrade, and Downgrade Instructions on page 23
- Product Compatibility on page 24

Known Behavior

There are no known limitations for the EX Series switches in Junos OS Release 16.2.

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 17
- Changes in Behavior and Syntax on page 17
- Known Issues on page 19
- Resolved Issues on page 20
- Documentation Updates on page 22
- Migration, Upgrade, and Downgrade Instructions on page 23
- Product Compatibility on page 24

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 16.2R2 for the EX Series.

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

- High Availability (HA) and Resiliency on page 19
- Network Management on page 19
- Platform and Infrastructure on page 20
- Port Security on page 20

High Availability (HA) and Resiliency

- On EX9200 Virtual Chassis, in a rare scenario, GRES might not reach the ready state and might fail to start, because the Routing Engine does not receive the state ack message from the Packet Forwarding Engine after performing GRES. This is a timing issue. It might also stop Routing Engine resource releasing and then cause resource exhausting. Reboot the system if this problem occurs. PR1236882

Network Management

- SNMP queries to retrieve jnxRpmResSumPercentLost will return the RPM/TWAMP probe loss percentage as an integer value whereas the precise value (including decimal points) can be retrieved through the CLI by using the following commands:
  - **show services rpm probe-results**
  - **show services rpm twamp client probe-results**

PR1104897

Platform and Infrastructure

- On EX9208 switches, a DCD restart might disable the member links in an MC-LAG, resulting in traffic loss. PR1229001

Port Security

- On an EX9200-6QS line card, storm control might not work for multicast traffic. PR1191611

See Also

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 19
- Resolved Issues on page 20
- Documentation Updates on page 22
- Migration, Upgrade, and Downgrade Instructions on page 23
- Product Compatibility on page 24

Resolved Issues

This section lists the issues fixed in the Junos OS main release and the maintenance releases for EX Series switches.

For the most complete and latest information about known Junos OS defects, use the Juniper online Junos Problem Report Search application.

- Resolved Issues: Release 16.2R2 on page 20

Resolved Issues: Release 16.2R2

- **Authentication and Access Control**
- **High Availability (HA) and Resiliency**
- **Interfaces and Chassis**
- **Port Security**
- **Routing Policy and Firewall Filters**
- **Routing Protocols**
- **Spanning-Tree Protocols**
- **VPNs**

Authentication and Access Control

- On EX9200 Virtual Chassis, MAC address learning might fail on an authenticated interface assigned to **voice-vlan** by dynamic VLAN assignment in single-secure mode. PR1212826
- 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. PR1233261

High Availability (HA) and Resiliency

- On EX9200 switches, if ISSU is used to upgrade Junos, it is possible that an unnecessary thread will run on an FPC after the upgrade procedure. This thread can potentially enter into a loop and trigger a stop of forwarding traffic on that particular FPC. PR1249375

Interfaces and Chassis

- On an EX9200 switch with MC-LAG, when the **enhanced-convergence** statement is enabled, and when the kernel sends a next-hop message to the Packet Forwarding Engine, the full Layer 2 header is not sent and a packet might be generated with an invalid source MAC address for some VLANs. PR1223662

Port Security

- 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
- A vulnerability in IPv6 processing has been discovered that may allow a specially crafted IPv6 Neighbor Discovery (ND) packet destined to an EX Series Ethernet Switches to cause a slow memory leak. A malicious network-based packet flood of these crafted IPv6 NDP packets may eventually lead to resource exhaustion and a denial of service. Use access lists or firewall filters to limit ICMPv6 traffic destined to the switch only from trusted hosts. PR1206593
- On EX9200 switches, after a MACSec link flaps, traffic is not forwarded across the MACSec link. PR1269229

Routing Policy and Firewall Filters

- 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

Routing Protocols

- On EX9200 switches, in a rare condition after a BGP session flaps, BGP updates might not be sent completely, resulting in BGP routes being shown in the advertising-protocol table on the local end but not shown in the receive-protocol table on the remote end. PR1231707

Spanning-Tree Protocols

- On EX9200 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
- On EX9200, the command **set protocols rstp interface all edge** configures all interfaces to go into BPDU block even if an interface is explicitly disabled under the **rstp** hierarchy. PR1266035

VPNs

- 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

See Also

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 19
- Documentation Updates on page 22
- Migration, Upgrade, and Downgrade Instructions on page 23
- Product Compatibility on page 24

Documentation Updates

There are no errata or changes in Junos OS Release 16.2R2 for the EX Series switches documentation.

See Also

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 19
- Resolved Issues on page 20
- Migration, Upgrade, and Downgrade Instructions on page 23
- Product Compatibility on page 24

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 23

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

See Also

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 19

Product Compatibility

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/.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the Hardware Compatibility tool.

See Also

- New and Changed Features on page 17
- Changes in Behavior and Syntax on page 17
- Known Behavior on page 18
- Known Issues on page 19
- Resolved Issues on page 20
- Documentation Updates on page 22
- Migration, Upgrade, and Downgrade Instructions on page 23

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

These release notes accompany Junos OS Release 16.2R2 for the 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/software/junos/.

New and Changed Features

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for MX Series and T Series.

Release 16.2R2 New and Changed Features

Interfaces and Chassis

- **Enhancement to ambient-temperature statement (MX Series)**—Starting in Junos OS Release 16.2R2 the default ambient temperature is set at 40° C on MX480, MX960, MX2010, and MX2020 Universal Routing Platforms. You can override ambient temperature by resetting the temperature to 55° C or 25° C.

```
[edit]
user@router# set chassis ambient-temperature ?
Possible completions:
25C                  25 degree celsius
40C                  40 degree celsius
55C                  55 degree celsius
[edit]
```

  When a router restarts, the system adjusts the power allocation or the provisioned power for the line cards on the basis of the configured ambient temperature. If there is not enough power, then a minor chassis alarm is raised. However, the chassis continues to run with the configured ambient temperature. You can configure a new higher ambient temperature only after you make more power available by adding new

power supply modules or by taking a few line cards offline. By using the provisioned power that is saved by configuring a lower ambient temperature, you can bring more hardware components online.

- **Enhancement to policer configuration**—Starting in Junos OS Release 16.2R2, you can configure the MPC to take a value from 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.

Subscriber Management and Services

- **LAC support for IPv6 address family and firewalls (MX Series)**—Starting in Junos OS Release 16.2R2, you can configure the LAC to create the IPv6 address family (inet6) when establishing a tunnel for subscribers. By default, the LAC requires only family inet to enable forwarding into an IP tunnel. It can apply IPv4 firewall filters to the session. Even when family inet6 is included in the dynamic profile, by default it is not created and IPv6 firewall filters cannot be applied.

  Include the **enable-ipv6-services-for-lac** statement at the **[edit system services l2tp]** hierarchy level to allow the IPv6 family to be created and IPv6 filters to be applied.

  Use the **show services l2tp summary** command to display the current state, Disabled or Enabled, in the IPv6 Services for LAC Sessions field.

- **Configurable grace period for unresponsive RADIUS servers (MX Series)**—Starting in Junos OS Release 16.2R2, 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 from 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 was 10 seconds and was not configurable.

- **Support for excluding tunnel attributes from RADIUS Access-Request messages (MX Series)**—Starting in Junos OS Release 16.2R2, you can use the **exclude** statement at the **[edit access profile *profile-name* radius attribute]** hierarchy level to exclude the following tunnel attributes from RADIUS Access-Request messages in addition to the previously supported Accounting-Start and Accounting-Stop messages:

  - acct-tunnel-connection—RADIUS attribute 68, Acct-Tunnel-Connection
  - tunnel-assignment-id—RADIUS attribute 82, Tunnel-Assignment-Id
  - tunnel-client-auth-id—RADIUS attribute 90, Tunnel-Client-Auth-Id

- tunnel-client-endpoint—RADIUS attribute 66, Tunnel-Client-Endpoint
- tunnel-medium-type—RADIUS attribute 65, Tunnel-Medium-Type
- tunnel-server-auth-id—RADIUS attribute 91, Tunnel-Server-Auth-Id
- tunnel-server-endpoint—RADIUS attribute 67, Tunnel-Server-Endpoint
- tunnel-type—RADIUS attribute 64, Tunnel-Type

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (MX Series and T Series)**—Starting in Junos OS Release 16.2R2, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database. The ephemeral database provides a fast programmatic interface that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. Junos OS provides a default instance and up to eight user-defined instances of the ephemeral configuration database. The device's active configuration is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database. Ephemeral configuration data is volatile and is deleted upon rebooting the device.

  [See Understanding the Ephemeral Configuration Database.]

Release 16.2R1 New and Changed Features

Authentication, Authorization and Accounting

- **Hardened shared secrets in Junos OS (MX Series)**—Starting in Junos OS Release 16.2, new CLI commands are introduced to configure a system master password and request to decrypt an encrypted secret, allowing for hardening of shared secrets such as preshared keys and RADIUS passwords.

  Having a master password enables devices to encrypt passwords in such a way that only devices running Junos OS that have knowledge of the master password can decrypt the encrypted passwords. The following new CLI commands are available:

  - **request system decrypt password**

- **set system master-password**

Class of Service

- **Support for ingress rate limiting (MX Series)**—Beginning with Junos OS Release 16.2R1, on MPCs that support ingress queuing, you can perform rate limiting on incoming packets based on the forwarding class and packet loss priority defined for each packet at ingress. You can perform ingress rate limiting by applying an input traffic control profile (using the **input-traffic-control-profile** statement) or an input scheduler map (using the **input-scheduler-map** statement) to a physical or logical interface where the traffic control profile or scheduler map contains a rate-limited scheduler.

  [See Ingress Rate Limiting on MX Series Routers with MPCs.]

EVPNs

- **EVPN MAC pinning (MX Series)**—Starting in Release 16.2, Junos OS enables MAC pinning for Ethernet VPN (EVPN), including customer edge (CE) interfaces and EVPN over MPLS core in both all-active mode and single-active mode.

  A MAC address pinned over CE interfaces in EVPN is synchronized to remote EVPN PE devices by adding the Sticky bit (in accordance with RFC 7432, Section 7.7, MAC Mobility Extended Community). On a remote EVPN PE device, a MAC address received with Sticky bit enabled is pinned over the MPLS core. A pinned MAC address cannot be moved to a different interface. When a MAC address is pinned locally in a bridge domain, the address is synchronized to remote EVPN PE devices. The interface with the locally pinned MAC address discards traffic sent from any other interface that has the identical MAC address if it is learned locally in the bridge domain.

- **Distribution of VXLAN VNIDs using EVPN (MX Series)**—Starting in Release 16.2, Junos OS enables Ethernet VPN (EVPN) with Virtual Extensible LAN (VXLAN) encapsulation to provide Layer 2 connectivity for endpoints within a virtual network that Contrail virtualization software creates. Endpoints in this scheme include virtual machines (VMs) connected to a virtual server, and nonvirtual bare-metal servers (BMSs) connected to a top-of-rack (ToR) platform. An MX Series router functions as a default gateway for nonvirtual BMSs for the traffic among the endpoints that belong to different virtual networks.

  The virtual network uses two types of encapsulation:

  - MPLS-over-GRE is used for L3 routing between Contrail and MX platform.
  - EVPN with VXLAN encapsulation is used for L2 connectivity between VM and BMS within a VN.

  An MX Series router supports all-active L3 gateways for redundancy and load balancing to ensure failure protection for the default gateway.

General Routing

- **Support for the combined operation of Synchronous Ethernet and Precision Time Protocol or hybrid mode (MX104)**—A combined operation of Synchronous Ethernet and Precision Time Protocol (PTP), also known as hybrid mode, is supported on the MX104 routers. In hybrid mode, the Synchronous Ethernet equipment clock (EEC) derives the frequency from Synchronous Ethernet and the phase and time of day from PTP (also known as IEEE 1588v2) for time synchronization.

  Synchronous Ethernet and PTP provide frequency and phase synchronization; however, the accuracy in the order of nanoseconds is difficult to achieve through either PTP or Synchronous Ethernet, and they do not support a large number of network hops. Hybrid mode resolves these issues by extending the number of network hops and also provides the clock synchronization accuracy in the order of tens of nanoseconds.

  

  NOTE: Hybrid mode is not supported on integrated routing and bridging (IRB) and aggregated Ethernet interfaces configured on MX104 routers.

High Availability and Resiliency

- **NSR support for EVPN (MX Series)**—Starting in Release 16.2, Junos OS ensures minimal loss of traffic when a Routing Engine switchover occurs with nonstop active routing (NSR) and graceful Routing Engine switchover (GRES) enabled. The forwarding state of the Packet Forwarding Engine (PFE) remains intact during switchover. The signaling state on the primary Routing Engine and on the standby Routing Engine are built in parallel.

  This feature is supported for EVPN over MPLS.

  

  NOTE: Expect a traffic loss pertaining to a topology change if the topology change occurs during a switchover.

- **PIM NSR support for VXLAN (MX Series)**—Starting in Release 16.2, Junos OS enables Protocol Independent Multicast (PIM) nonstop routing (NSR) support for Virtual Extensible LANs (VXLANs).

  The Layer 2 address learning process (l2ald) passes VXLAN parameters (vxlan multicast group addresses and vtep-interface-source) to the routing protocol process on the master Routing Engine. The routing protocol process forms PIM joins with the multicast routes through the pseudo-VXLAN interface.

  Because the l2ald does not run on the backup Routing Engine, the PIM NSR mirroring mechanism provides the VXLAN configuration details to the backup Routing Engine. The routing protocol process matches the multicast routes on the backup Routing Engine with PIM states, which maintains the multicast routes in the Forwarding state.

- Unified ISSU support (MX104)—Unified in-service software upgrade (ISSU) is supported on the MX104 router. 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. [See Unified ISSU Concepts]

- **MX Series Virtual Chassis ISSU support for MPC6E line cards (MX Series Virtual Chassis)**—Starting in Junos OS Release 16.1, MPC6E line cards support ISSU in MX Series Virtual Chassis environments.

Interfaces and Chassis

- **Support for Synchronous Ethernet and Synchronization Status Messages on MIC-3D-4OC3OC12-1OC48 and MIC-3D-16CHE1-T1-CE (MX104, MX240, MX480, MX960)**—Starting with Junos OS Release 16.2R1, Synchronous Ethernet and Synchronization Status Messages (SSMs) are supported on the MIC-3D-4OC3OC12-1OC48 and MIC-3D-16CHE1-T1-CE MICs. Synchronous Ethernet (ITU-T G.8261 and ITU-T G.8264) is a physical layer technology that enables you to deliver synchronization services. It supports sourcing and transfer of frequency for synchronization purposes for both wireless and wireline services. An SSM indicates the quality level of the transmitting synchronous Ethernet Equipment Clock (EEC).

  You can configure Channelized T1 (**ct1**) interfaces as clock sources on MIC-3D-4OC3OC12-1OC48 and MIC-3D-16CHE1-T1-CE. To configure a clock source, you must specify the parameters that must be considered by the clock selection algorithm while selecting the best clock source. The parameters include the quality level value, the priority of the clock source, the request criteria, and the wait time to restore the interface signal to up state. To configure **ct1** as a clock source, include the **set source interfaces** *interface-name* statement at the **edit [chassis synchronization]** hierarchy level.

  NOTE:
  - To configure the ct1 interface as a clock source, ensure that the option-2 network option is configured.
  - You can configure a maximum of eight clock sources by using the set chassis synchronization source *source-name* command. If you attempt to configure more than eight sources, the configuration fails.
  - You can configure the ct1 interface to enable Ethernet Synchronization Message Channel (ESMC) packet transmission by using the set chassis synchronization esmc-transmit interfaces *interface-name* command.

IPv6

- **Forced IPv6 DNS server address insertion (MX Series)**—Starting in Junos OS Release 16.2, MX Series devices can dynamically provision DHCPv6 lease times and DNSv6 Server IP addresses for DHCPv6 clients. The IP addresses and lease times are provided to DHCPv6 clients in DHCPv6 Advertisement and Reply messages without requiring a Solicit or Request message from a CPE device.

Layer 2 Features

- **Implicit maximum bandwidth for inline services for L2TP LNS (MX Series)**—Starting in Junos OS Release 16.2, 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

- Starting in Junos OS Release 16.2R1, a new framework for API clients that uses the gRPC protocol is available for session management and device interaction. The gRPC protocol provides the request/response interface between the Junos extension toolkit (JET) service daemon (JSD) and the on-box or off-box application. The gRPC framework replaces the Apache Thrift framework that was used in previous releases. [See www.grpc.io.]

- **New Programmable Routing Protocol Process (prpd) Configuration Statements and Operational Commands (MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series)**—Starting in Junos OS Release 16.2R1, new configuration statements are introduced to allow you to set purge a timeout for prpd API clients and to set traceoptions to log information regarding those clients. A new operational command is introduced to allow you to monitor prpd clients and information regarding their connections.

  [See set routing-options programmable-rpd purge-timeout, set routing-options programmable-rpd traceoptions flag]

- **Support for adding nonnative YANG RPCs to the Junos OS schema (MX Series and T Series)**—Starting with Junos OS Release 16.1R3, you can load custom YANG RPCs on devices running Junos OS. Creating custom RPCs enables you to precisely define the input parameters and operations and the output fields and formatting for your specific operational tasks on those devices. The ability to add custom RPCs to a device is also beneficial when you want to create RPCs that are device-agnostic and vendor-neutral. You can load YANG modules that add custom RPCs by using the **request system yang add** operational command.

  [See Creating Custom RPCs in YANG for Devices Running Junos OS.]

Network Management and Monitoring

- **SNMP support for the timing feature on MX104 routers**—Starting in Junos OS Release 16.2R1, SNMP supports the timing feature on MX104 routers. Currently, SNMP support is limited to defect and event notifications through SNMP traps. The enterprise-specific MIB, Timing Feature Defect/Event Notification MIB, helps to monitor the operation of PTP clocks within the network. The trap notifications are disabled by default. To enable trap notifications for the timing events and defects , include the **timing-event** statement at the **[edit snmp trap-group *trap-group* object categories]** hierarchy level.

Platform and Infrastructure

- **Virtual broadband network gateway support on virtual MX Series router (vMX)**—Starting in Junos OS Release 16.2, vMX supports most of the subscriber management features available with Junos OS Release 16.2 on MX Series routers to provide a virtual broadband network gateway on x86 servers.

  vBNG runs on vMX, so it has similar exceptions; the following subscriber management features available on MX Series routers are not supported for vBNG:

  - High availability features such as hot-standby backup for enhanced subscriber management and MX Series Virtual Chassis.
  - CoS features such as shaping applied to an agent circuit identifier (ACI) interface set and its members.

  To deploy a vBNG instance, you must purchase these licenses:

  - vMX PREMIUM application package license with 1 Gbps, 5 Gbps, 10 Gbps, or 40 Gbps bandwidth
  - vBNG subscriber scale license with 1000, 10 thousand, 100 thousand, or 1 million subscriber sessions for one of these tiers: Introductory, Preferred, or Elite

- **Virtual MX Series router (vMX)**—Starting in Junos OS Release 16.2, you can deploy vMX routers on x86 servers. FreeBSD 10 is the underlying OS for Junos OS for vMX.

  vMX supports most of the features available on MX Series routers and allows you to leverage Junos OS to provide a quick and flexible deployment. vMX provides the following benefits:

  - Optimizes carrier-grade routing for the x86 environment
  - Simplifies operations by consistency with MX Series routers
  - Introduces new services without reconfiguration of current infrastructure

Routing Protocols

- **Support for OSPF segment routing (MX Series)**—Starting with Junos OS Release 16.2R1, IPv4 OSPF segment routing support is enabled through MPLS. OSPF creates an adjacency segment per OSPF neighbor, for a given interface, adjacency, and area. A separate MPLS label is allocated for each adjacency segment created.

Labels are allocated only when the neighbor moves from **Init** state to **Up**state and requests the label manager for an unreserved label. The corresponding label transitions are downloaded to the MPLS forwarding table after the label is advertised in locally originated LSPs. In case of LAN adjacencies, OSPF neighborship remains in a two-way state for the adjacencies between the DR-others. A separate label is allocated for each of the LAN neighbors, including the DR-other adjacencies that remain in the two-way state.

The Junos OSPF implementation enables the network operator to provision the following:

- IPv4 address family node segment index **node-sid**—This node-sid will be assigned to a router and used by all other remote routers in the network to index into respective node segment label blocks (SRGBs). It derives the segment identifier to forward IPv4 traffic destined for the same router which was assigned as node-sid.

NOTE: **Provisioning the IPv4 node-sid is allowed per routing instance, and is not allowed per OSPF area.**

[See Source Packet Routing in Networking (SPRING).]

- **Support for dynamic GRE tunnel creation based on IPv6 and 6VPE routes (MX Series)**— Starting with Junos OS Release 16.2, dynamic GRE tunnel creation is triggered by IPv6 L3VPN as well as 6VPE routes without the preexistence of IPv4 L3VPN routes in the same VRF instance.
- **Support of inner-vlan-list for qualified-bum-pruning-mode of VPLS routing instance (MX Series with MPCs/MICs FPCs)**— Starting with Junos OS Release 16.2, support for **qualified-bum-pruning-mode** is provided on dual-tagged subscriber interfaces configured with inner VLAN list or inner VLAN range for a VPLS routing instance. This allows the BUM traffic egressing this interface to be checked against the combination of the single service provider VLAN with the inner VLAN list or inner VLAN range of the subscriber interface and forward only the packet that is intended for the subscriber. The inner VLAN list on a subscriber interface can have multiple elements. Each element of the inner VLAN list can be a single VLAN tag or a range of VLANs.
- **Enhancement to the output of the show route detail operational command (MX80, MX104, MX240, MX480, MX960, MX2010, MX2020, vMX Series)**—Starting in Junos OS Release 16.2, the output of the **show route detail** command has been enhanced to show the keyword **programmed** in the state output field if the route was installed programmatically by an API client application.

  [See show route detail.]
- **BGP labeled unicast supports stack of labels (MX Series)**—Beginning with Release 15.1F5, 16.2R1, and later releases, Junos OS supports RFC 3107, *Carrying Label Information in BGP-4*, that allows stacking of multiple labels in the BGP labeled unicast. In earlier Junos OS Releases, only one label per prefix is supported in the BGP unicast label. Junos OS now supports a label stack of up to five labels per prefix in the BGP labeled unicast updates. BGP labeled unicast updates with more than five labels are not supported and Junos OS sets their state to **hidden**. This feature allows the use of BGP

unicast label stack to control packet forwarding in the network and to reflect the BGP unicast label stack routes to its clients without changing the next hop.

- **Support for IS-IS flooding groups (MX Series and T Series)**—Starting with Junos OS Release 15.1F5, 16.2R1, and later releases, you can configure flooding groups with IS-IS. This feature limits link-state PDU flooding over IS-IS interfaces.

  An LSP that is not self-originated is flooded only through the interface belonging to the flood group that has the configured area ID in the LSP. This helps minimize the routes and topology information, thus ensuring optimal convergence. You can segregate both level 1 and level 2 networks into flood groups by using area IDs as tags to identify a flood group. Configure interfaces with specific area IDs to modify the flooding behavior as per your requirements.

  To enable IS-IS flooding groups, include the **flood-group *flood-group-area-ID*** statement at the **[edit protocols isis interface]** hierarchy level.

  [See IS-IS Overview.]

- **Micro loop avoidance when IS-IS link fails (MX Series and T Series)**—Beginning with Release 15.1F5, 16.2R1, and later releases, Junos OS enables a device to defer IS-IS route download when an IS-IS link fails in order to avoid micro loops. When local links go down, the IS-IS protocol floods an entire area with the database. If the node connected to the local interface that has failed converges faster than the neighboring node, then the connected node redirects traffic to the converged path. This redirection can result in micro looping of traffic until the neighboring node converges. When the primary path of a protected node fails, the connected node does not need to converge quickly if the configured backup path is not impacted. In this case, traffic flow toward a converged path is deferred until the configured delay time.

- **Support for IS-IS segment routing (MX Series)**—Starting with Junos OS Release 15.1F5, 16.2R1, and later releases, IS-IS segment routing support is enabled through MPLS. Currently, label advertisements are supported for IS-IS only. IS-IS creates an adjacency segment per adjacency, per level, and per address family (one each for IPv4 and IPv6). Junos OS IS-IS implementation allocates node segment label blocks in accordance with the IS-IS protocol extensions for supporting segment routing node segments and provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:

  - **source-packet-routing**—Enable the source packet routing feature.
  - **node-segment**—Enable source packet routing at all levels.
  - **use-source-packet-routing**—Enable use of source packet routing node segment labels for computing backup paths for normal IPv4 or IPv6 IS-IS prefixes and primary IS-IS source packet routing node segments.
  - **no-advertise-adjacency-segment**—Disable advertising of the adjacency segment on all levels for a specific interface.

  [See Source Packet Routing in Networking (SPRING).]

- **Support for BGP Optimal Route Reflection (BGP-ORR) (MX Series)**—Starting with Junos OS Release 16.2, you can configure BGP-ORR with OSPF as the interior gateway

protocol (IGP) on a route reflector to advertise the best path to the BGP-ORR client groups by using the shortest IGP metric from a client's perspective, instead of the route reflector's view.

To enable BGP-ORR, include the **optimal-route-reflection** statement at the **[edit protocols bgp group *group-name*]** hierarchy level.

Use the following CLI commands to monitor and troubleshoot the configuration for BGP-ORR:

- **show bgp group**—View the primary and backup configurations of BGP-ORR.
- **show ospf bgp-orr**—View the OSPF BGP-ORR metric (RIB).
- **show route advertising protocol bgp *peer***—Verify whether the routes are being advertised according to the BGP-ORR rules.

Security

- **Global configuration for flow detection and tracking (MX Series)**—Starting in Junos OS Release 16.2, you can configure the mode of operation for flow detection and tracking globally for all protocol groups and packet types. In earlier releases, although you enable flow detection and tracking globally, you can configure the behavior only at the individual flow aggregation levels: physical interface, logical interface, or subscriber; you cannot configure the behavior globally. The new global configuration applies to all packet types in the traffic flow unless it is overridden by the configuration for a protocol group or packet type at the flow aggregation levels.

Services Applications

- **Support for load balancing dynamic endpoint IPsec tunnels among services interfaces (MX Series routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 16.2, you can load-balance IPsec tunnels with dynamic endpoints among services interfaces on multiple PICs. You configure load balancing by configuring multiple next-hop IPsec service sets, each identifying services interfaces with the **inside-service** and **outside-service** statements at the **[edit services service-set *service-set-name* next-hop-service]** hierarchy level. The services interfaces in the **outside-service** statements of the service sets must be in the same VPN routing and forwarding (VRF) instance, and each of the service sets must include the same **local-gateway** and **ike-access-profile** values at the **[edit services service-set *service-set-name* ipsec-vpn-options]** hierarchy level.

  [See Configuring Dynamic Endpoints for IPsec Tunnels.]

- **Support for multiple source-destination port pairs in a DTCP ADD request (MX Series routers)**—Starting in Junos OS Release 16.2, the MX Series router can process mediation device DTCP ADD requests that contain up to 15 source-destination port pairs. Multiple source-destination port pairs must be separated by commas.

Subscriber Management and Services

- **Shared memory logs for client-level analysis (MX Series)**—Starting in Junos OS Release 16.2, shared memory logging **smlog** is available by default so you can retrieve client activity on a per subscriber basis. Additional filters enables you to retrieve logs according to a variety of parameters, including the client identifier, client DUID, interface name, IP address, session ID, subnet, and VLAN.

  A complete list of supported filters is available at the **[show shmlog entries logname all]** hierarchy level and a complete list of flags is available at the **[show shmlog entries logname all flag-name]** hierarchy level.

  See *shmlog* for information about disabling shared memory logs, and steps required to view the data.

- **Support dynamic VLAN access profile assignments (MX Series)**—Starting in Junos OS Release 16.2, you can assign different access profiles to different dynamic profiles on the same interface. For example, you can attach an access profile to an interface configured for dynamic VLAN/SVLAN so all the VLANs/SVLANs use the same set of authentication, authorization, and accounting parameters. However, different access profiles can have different authentication/authorization settings, so that you could have authentication on some VLAN/SVLAN ranges but not on other ranges.

  To configure dynamic VLAN access profile assignments, add the access profile at the **[edit interfaces <interface-name> auto-configure vlan-ranges dynamic-profile <profile-name> access-profile <vlan-access-profile-name>]** or **[edit interfaces <interface-name> auto-configure stacked-vlan-ranges dynamic-profile <svlan-profile-name> access-profile <svlan-vlan-access-profile-name>]** hierarchy level.

- **Support for 1:1 LNS stateful redundancy on aggregated inline service interfaces (MX Series with MPCs and MIC interfaces)**—Starting in Junos OS Release 16.2, you can create an aggregated inline service virtual logical interface that bundles pairs of inline services anchor interfaces across MPCs to provide 1:1 LNS stateful redundancy between the paired members. You can assign a single bundle or one or more pools of bundles per L2TP tunnel group.

  LNS sessions are subsequently established on the aggregated interface. When an LNS session failover occurs, the secondary link becomes active and all the LNS data traffic destined for the session automatically moves over to the secondary anchor interface on a different MPC. The subscriber session remains up on the virtual logical interface. No traffic statistics are lost. If this redundancy is not configured, subscriber traffic is lost, the keepalives expire, and the PPP client is disconnected.

  When a card comes back online after a failover, you can move the LNS data traffic from the currently active secondary interface back to the primary interface on the original card You can manually force a switchover from the primary interface to the secondary interface, and you can manually revert to the original interface in this case as well.

- **Subscriber login session with optional services (MX Series)**—Starting in Junos OS Release 16.2, you can use the **service activation** statement at the **[edit access profile *profile-name* radius options**] hierarchy level to specify whether successful activation

of services referenced in the Activate-Service VSA (26-65) in the RADIUS Access-Accept message is required or optional for subscriber login access.

When activation is required, failure for any reason causes the Network-Family-Activate-Request for that network family to fail. If no other network family is already active for the subscriber, then the client application logs out the subscriber.

When activation is optional, subscribers can still log in when a service fails to activate because of a configuration error. Failures for any other reason do not allow successful login.

By default, activation is required for services applied with a dynamic profile and is optional for services applied by an Extensible Subscriber Services Manager (ESSM) operation script. In earlier releases, only the default behavior is available.

NOTE: This configuration does not apply to services activated by means of RADIUS CoA requests, JSRC Push-Profile-Request (PPR) messages, or subscriber secure policy.

- **Processing multiple activation and deactivation requests in a single CoA message (MX Series)**—Starting in Junos OS Release 16.2R1, subscriber management processes RADIUS-initiated Change of Authorization (CoA) messages in a more efficient manner. When it receives CoA message that has multiple activation and deactivation requests, the router groups the requests together, by type. The router then processes all deactivation requests before processing the activation requests.

  Processing deactivation requests first helps the router provide a consistent behavior for activated services. For example, a particular service might be activated multiple times, using different parameters. It is more efficient for the router to process the deactivation requests for existing instances of the service before attempting to activate the same service with different parameters.

  In earlier releases, the router processed all activation requests first, before processing the deactivation requests in the CoA message.

- **Support for username stripping per routing instance (MX Series)**—Starting in Junos OS Release 16.2, you can configure a subscriber access profile so that a portion of each subscriber login string is discarded and the remaining characters subsequently are used as a modified username by an external AAA server for session authentication and accounting. The modified username appears, for example, in RADIUS Access-Request, Acct-Start, and Acct-Stop messages, as well as RADIUS-initiated disconnect requests and change of authorization (CoA) requests. The login string is examined in the direction you specify until a delimiter is identified as one of the configured delimiters. The delimiter and all characters to the right of the delimiter are discarded.

- **AAA option sets to authorize and configure subscribers per routing instance to support username stripping (MX Series)**—Starting in Junos OS Release 16.2, you can include one or more of the following statements at the new **[edit access aaa-options *aaa-options-name*]** hierarchy level to define a set of AAA options for a subscriber or set of subscribers that username stripping is applied to:

- **access-profile** *profile-name*—Specify the name of the access profile that includes the username stripping configuration.
- **aaa-context** *aaa-context-name*—Specify the logical-system:routing-instance that the subscriber session uses for AAA (RADIUS) interactions like authenticating and accounting.
- **subscriber-context** *subscriber-context-name*—Specify the logical-system:routing-instance in which the subscriber interface is placed.

NOTE: Only the default (master) logical system is supported.

Use the **aaa-options** *aaa-options-name* statement at the **[edit dynamic-profiles** *profile-name* **interfaces pp0 unit $junos-interface-unit ppp-options]** hierarchy level to apply the attributes to PPP subscribers tunneled from the LAC to the LNS inline service interface.

Alternatively, use the **aaa-options** *aaa-options-name* statement at the **[edit access group-profile** *profile-name* **ppp-options]** hierarchy level to apply the attributes to PPP subscribers tunneled from LACs that are members of the user group.

Usernames are examined and modified according to the subscriber and AAA contexts specified in the option set. In the event of a conflict between option sets configured in both a group profile and a dynamic profile, the dynamic profile takes precedence.

- **Support for maximum session limits on L2TP service interfaces (MX Series)**—Starting in Junos OS Release 16.2, you can include the **l2tp-maximum-session** *number* statement at the **[edit interfaces si***service-interface***]** or **[edit interfaces asi***service-interface***]** hierarchy level to specify the maximum number of sessions that are allowed on an individual service interface (si) or aggregated service interface (asi). New session requests on an interface are accepted only when the session count is less than the maximum session limit. If the limit has been reached, subsequent requests are dropped and the LNS responds with a CDN message (Result Code 2, Error Code 4). When a pool of interfaces is configured, interfaces at the maximum limit are ignored in favor of an interface in the pool that has a lower session count. For an asi interface, the configuration applies to all member interfaces; you cannot configure the limit for individual member interfaces.
- **Enhanced load balancing on L2TP physical service interfaces (MX Series)**—Starting in Junos OS Release 16.2, when a service interface in a service device pool is rebooted, sessions reconnect and new session requests are distributed based on the number of sessions on the available interfaces in the pool. The sessions are assigned to the interface with the fewest sessions. If more than one interface has the minimum number of sessions, then a random selection determines which interface gets the session.

  In earlier releases, session load balancing is a simple round-robin distribution among the interfaces. Consequently, fewer sessions are assigned to a newly rebooted interface than to the other interfaces. For example, consider a pool with two si interfaces, si-0/0/0 and si-1/0/0. Each has 100 sessions. If si-1/0/0 reboots, it drops all 100 sessions. As the sessions reconnect, they alternate between the two interfaces so that

when all sessions have reconnected, si-0/0/0 has 150 sessions and the reconnected si-1/0/0 interface has only 50 sessions.

Consider the same pool with the new behavior. As sessions reconnect, si-1/0/0 has fewer sessions (0 to start) than si-0/0/0 (100). Because the interface with the fewest sessions is selected, all sessions are assigned to si-1/0/0 until it reaches the same count as si-0/0/0.

For aggregated services interfaces (asi), the interface with the lowest session count is selected from the pool for new or reconnect session requests. When the active si interface in the asi bundle goes down, all the active sessions on that primary interface fail over to the secondary interface.

- **Monitoring only ingress traffic for subscriber idle timeouts (MX Series)**—Starting in Junos OS Release 16.2, you can specify that only ingress data traffic is monitored for subscriber idle timeout processing. If you include the **client-idle-timeout-ingress-only** statement in addition to the **client-idle-timeout** statement at the **[edit access-profile *profile-name* session-options]** hierarchy level, subscribers are logged out or disconnected when no ingress traffic is received for the duration of the idle timeout period. Egress traffic is not monitored. If you do not include the **client-idle-timeout-ingress-only** statement, both ingress and egress data traffic are monitored during the timeout period to determine whether subscribers are logged out or disconnected.

- **Broadband-specific support for PCEF (MX Series)**—Starting in Junos OS Release 16.2, the policy and charging enforcement function (PCEF) is supported for broadband-specific functionality. PCEF is one of the major components of the 3rd Generation Partnership Project (3GPP) policy and charging control (PCC) architecture that provides the unification of wireline provisioning and accounting for customers. PCEF provides user traffic handling and CoS at the gateway, provides service data flow detection, and applies the rules received from the Policy and Charging Rules Function (PCRF). PCEF optionally interacts with the Online Charging System (OCS) using the 3GPP Gy protocol to retrieve policy and charging authorization for quotas and credit control.

  BPCEF configuration consists of configuring the characteristics of the PCRF and OCS with which it interacts. Include the **pcrf** statement at the **[edit access]** hierarchy level to configure the PCRF partition and the global attributes, rules, and parameters to authorize and provision subscribers. Include the **ocs** statement at the **[edit access]** hierarchy level to configure the OCS global attributes and partition. Include the **provisioning-order pcrf** statement at the **[edit access profile *profile-name*]** hierarchy level to request provisioning from the PCRF over the Gx protocol. Include the **charging-service-list ocs** statement at the **[edit access profile *profile-name*]** hierarchy level to configure the list of charging services to be communicated.

  Use these new commands to display PCRF and OCS information: **show network-access pcrf state**, **show network-access pcrf statistics**, **show network-access ocs state**, and **show network-access ocs statistics**. Use these new commands to clear statistics and subscriber counters: **clear network-access pcrf statistics**, **clear network-access pcrf subscribers**, and **clear network-access ocs statistics**.

- **DHCPv6 relay agent supports multiple addresses or prefixes per DUID (MX Series)**—Starting in Junos OS Release 16.2, DHCPv6 relay agent supports multiple

address or network prefix leases assigned to a single DHCP Unique ID (DUID). Existing operational commands that display DHCPv6 relay bindings now display multiple addresses and network prefixes. When you are configuring DHCPv6 relay agent, if service accounting is required separately for each address or network prefix issued to a single subscriber, you must configure a separate address pool at the DHCPv6 server for each address or network prefix allocated.

- **Support for vendor-specific information in DHCPv4 and DHCPv6 relay (MX Series)**—Starting in Junos OS Release 16.2, you can add a hostname, location (such as a unique connection identifier), or both in DHCP control packets sent over server-facing interfaces. For DHCPv4 relays, the feature leverages option 82, suboption 9, to provide the vendor-specific information. For DHCPv6 relays, it is added under the vendor-specific option (17).

  This feature can be useful when used with operator-developed tools for troubleshooting DHCP servers and providing service assurances. For example, a central DHCP server can log the information, and operators can query a single entity to track and troubleshoot subscriber IP information and network attachment points.

  To configure vendor-specific information, add a hostname, location, or both at the **[edit forwarding-options dhcp-relay relay-option-82 vendor-specific host-name]** or **[edit forwarding-options dhcp-relay relay-option-82 vendor-specific location]** hierarchy level.

- **Preserving and restoring IPv6 prefixes assigned using DHCPv6 Prefix Delegation (MX Series)**—Starting in Junos OS Release 16.2, when IPv6 addresses are assigned using DHCPv6 Prefix Delegation, you can configure the router to preserve and restore a subscriber's delegated prefix through multiple logins. This feature prevents an IA-PD change, which triggers renegotiation for all hosts attached to the residential gateway. This feature requires the use of Agent-Circuit-IDs (ACIs) to identify subscribers.

- **Subscriber management and services feature and scaling parity (MX104)**—The MX104 router supports all subscriber management and services features that are supported by the MX80 router and the MX240, MX480, and MX960 routers. In addition, the scaling and performance values for the MX104 router match those of the MX80, MX240, MX480, and MX960 routers.

- **DHCP rate adjustment (MX Series)**—Starting in Junos OS Release 16.2, 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.

System Logging

- **Internal communication health monitor**—Starting in Junos OS Release 16.2, you can use the internal communication health monitor daemon (icmd). The icmd daemon runs on all Routing Engines in the chassis or virtual chassis and it monitors internal communication systems, reports any unexpected communication changes, logs communication issues to **/var/log/messages** and provides debugging information.

  The icmd daemon can be deployed on all platforms as **icmd** in **/usr/libexec** and runs automatically with jlaunchd.

VPNS

- **Support for next-hop-based dynamic tunnels (MX Series and T Series)**—Starting with Junos OS Release 16.2, dynamic generic routing encapsulation (GRE) and UDP tunnels support the creation of a tunnel composite next hop for every dynamic tunnel created. The tunnel composite next hop includes the dynamic tunnel's source and destination IP address, encapsulation data, and a VPN label (when chained composite next hop is not enabled).

  For dynamic GRE tunnels, the next-hop-based tunneling feature should be explicitly configured. This feature overcomes the scaling limitation of the default interface-based tunnel mode by removing the dependency on physical interfaces, and creating a next-hop ID instead of a next-hop interface for every GRE tunnel configured. To enable next-hop-based dynamic GRE tunnels, include the **next-hop-based-tunnel** statement at the **[edit routing-options dynamic-tunnels gre]** hierarchy level. With next-hop-based dynamic GRE tunnels, a device can scale up to 32,000 dynamic GRE tunnels.

  For dynamic UDP tunnels, the next-hop-based tunnel mode is supported by default. These tunnels are referred to as MPLS-over-UDP tunnels, and they provide a scaling advantage of up to 4,000 UDP tunnels on a device.

  The next-hop-based dynamic tunnel feature benefits data center deployments that require mesh IP connectivity from one provider edge (PE) device to all other PE devices in the network. At a given point in time, for the same tunnel destination, the next-hop-based dynamic tunnel encapsulation can either be GRE or UDP.

  [See Example: Configuring Next-Hop-Based MPLS-Over-UDP Dynamic Tunnels and Example: Configuring Next-Hop-Based Dynamic GRE Tunnels.]

VXLAN

- **Overlay ping and traceroute functionality for VXLAN tunnels (MX Series)**—Starting in Junos OS Release 16.2R1, two new CLI commands supporting ping and traceroute troubleshooting functionality are provided to debug VXLAN overlay tunnels: ***ping overlay vni vni-id tunnel-src ip-address-src tunnel-dst ip-address-dst*** and ***traceroute overlay vni vni-id tunnel-src ip-address-src tunnel-dst ip-address-dst***. Use the ping overlay and traceroute overlay commands to validate and verify the presence of the VXLAN tunnel endpoints within the context of the overlay VXLAN network identifier or VXLAN Segment ID (VNI) segment.Understanding Overlay ping and traceroute Packet Support

See Also
- Changes in Behavior and Syntax on page 43

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 16.2R2 for MX Series and T Series.

General Routing

- **Support for deletion of static routes when the BFD session goes down (MX Series)**—Starting with Junos OS 16.2R2, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

Interfaces and Chassis

- **Changes to show interfaces *interface-name* extensive output**—Starting in Junos OS Release 16.1R5 and 16.2R2, the MAC Control Frames field of the **show interface *interface-name* extensive** command for a specified 10-Gigabit Ethernet interface displays a value of zero. In previous releases, the value for this field was calculated. Because of continuous traffic and as a result of the calculations, the value displayed for this field changed continuously.

Junos OS XML API and Scripting

- **XML output change for show subscribers summary port command (MX Series)**—Starting in Junos OS Release 16.2R2, 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>
</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>
</rpc-reply>
```

Management

- **Support for status deprecated statement in YANG modules (MX Series and T Series)**—Starting with Junos OS Release 16.1R2, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.
- **XPath expressions for specific YANG keywords disabled during commit operations (MX Series and T Series)**—Starting in Junos OS Release 16.1R2, XPath expression evaluations for the following YANG keywords are disabled by default during commit operations: **leafref**, **must**, and **when**. Prior to Junos OS Release 16.1R2, Junos OS evaluates the constraints for these keywords, which can result in longer commit times.

MPLS

- **New field for LSP ping egress interface failure**—Starting in Junos OS Release 16.2R2, if an LSP ping is started and the chosen egress interface fails, pings are still sent to the failed interface and then dropped. The ping must be manually stopped and restarted to select a working interface to the destination (if one exists). To help detect this ping situation, a new field, **Packets dropped due to ifl down**, has been added to the output of the **show system statistics mpls** command.

  [See show system statistics mpls]

- **RSVP LSP Attribute Order Complies with RFC6510 (M Series, MX Series, and T Series)**—The Junos OS RSVP PATH/RESV messages follow the recommendations made in RFC6510 for the LSP attribute order.

Network Management and Monitoring

- **Possible change is in the object identifier (MX Series and T Series)**—The many warnings that occurred previously during MIB loading in the SNMP client has been reduced. This update could change the resulting OID depending on the SNMP client-loading logic. For example, jnxProductACX1000 is defined under the two following nodes:

  - jnxProductACX1000 OBJECT IDENTIFIER ::= { jnxProductVariationACX 1 }
  - jnxProductACX1000 OBJECT IDENTIFIER ::= { jnxProductVariationACX1000 1 }

  Because the second definition is the duplicate, it is removed. If previously, the SNMP client referred to the second OID based on MIB loading logic, then you would see a change in OID for the client.

- **Update to SNMP support of apply-path statement (MX Series)**—In Junos OS Release 16.2R2, SNMP implementation for the **apply-path** configuration statement supports only two lists:

  - **apply-path "policy-options prefix-list <list-name> <*>"**

    This configuration has been supported from day 1.

  - **apply-path "access radius-server <*>"**

This configuration is supported as of this release.

- **MIB buffer overruns can only be counted under ifOutDiscard (MX Series)**—Starting in Junos OS Release 16.2R2, qdrops (buffer overruns) are no longer counted under ifOutErrors along with ifOutDiscards. This contradicted RFC 2863, where buffer overruns should only be counted under ifOutDiscards and not under ifOutErrors.
- **SNMP syslog messages changed (MX Series)**—In Junos OS Release 16.2R2, two misleading SNMP syslog messages were rewritten to accurately describe the event:
  - OLD --**AgentX master agent failed to respond to ping. Attempting to re-register**
    NEW –- **AgentX master agent failed to respond to ping, triggering cleanup!**
  - OLD –- **NET-SNMP version %s AgentX subagent connected**
    NEW --- **NET-SNMP version %s AgentX subagent Open-Sent!**

  [See the MIB Explorer.]
- **Improved usage of wildcard in SNMP notify-filter OID (MX Series)**—In Junos OS Release 15.1R7, the filter subtree using an asterisk (*) is correctly read as a wildcard character and not as an ASCII value of 42. This issue previously occurred in the following routers:
  - M Series running Junos OS Release 11.4R13.5, 13.3R7-S1
  - ACX2000 Series running Junos OS Release 12.3X54-D20.9
  - MX Series running Junos OS Release 14.1X50-D125

  A sample of the change appears in the output of the **show snmp v3** command:

  Old Output

```
Filter              Subtree           Filter      Storage      Status
name                                  type        type
nf1                 1.2.3.4.5         include     nonvolatile  active
  <<<<< Issue
nf1                 1.42.6            include     nonvolatile  active
  <<<< Issue
```

  New Output

```
Filter              Subtree           Filter      Storage      Status
name                                  type        type
nf1                 1.*.*.4.5         include     nonvolatile  active
  <<< Fixed
nf1                 1.*.6             include     nonvolatile  active
   <<< Fixed
```

  [See the SNMP MIB Explorer.]
- **Enhancement to SMNPv3 traps for contextName field (MX Series)**—Starting in Junos OS Release 16.2R2 and Release 17.2, the contextName field in SNMPv3 traps generated from a non-default routing instance, is populated with the same routing-instance

information as is given in SNMPv2 traps. SNMPv2 traps provide the routing-instance information as context in the form of context@community. This information gives the network monitoring system (NMS) the origin of the trap, which is information it might need. But in SNMPv3, until now, the contextName field was empty. For traps originating from a default routing instance, this field is still empty, which now indicates that the origin of the trap is the default routing instance.

See SNMP MIB Explorer.

- **Juniper MIBs Loading Errors Fixed (MX Series)**—In junos OS Release 16.2R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:
  - jnx-gen-set.mib
  - jnx-ifotn.mib
  - jnx-optics.mib

  [See MIB Explorer.]

Operation, Administration, and Maintenance (OAM)

- **Change in behavior of the Ethernet OAM CFM process (MX Series)**—When you deactivate the connectivity fault management (CFM) protocol, the CFM process (cfmd) stops. When you activate the CFM protocol, cfmd starts.

  Prior to Junos OS Release 16.1R1, when you deactivate the CFM protocol, the CFM process continues to run.

- **Change in status of CFM sessions (MX Series with MPCs)**—Starting in Junos OS Release 16.2R1, for connectivity fault management (CFM) up MEP sessions on trunk interfaces, when the physical link is down and if the router's network service mode is configured as **enhanced-ip**, the local CFM session fails and the status of the CFM session displays **Failed**.

  In releases before Junos OS Release 16.2R1, when the physical link is down, the local CFM session does not fail and the status of the CFM session displays **OK**.

Platform and Infrastructure

- The length of TACACS messages allowed on Junos OS devices has been increased from 8150 to 65535 bytes. PR1147015
- **Improvements to MIB validation during Junos OS build (libjsnmp) (MX Series and T Series)**—Many warnings can be ignored if produced while MIBs are compiling. The following are some warnings should be considered errors because they can break the build:

  [0-9]:.*failed to locate—An OID that has failed to be located.

  [0-9]:.*redefinition of identifier—Redefinition of OIDs found in jnx-chas-defines.

  [0-9]:.*sequence-type-mismatch—Type mismatch found in sequence syntax of the table and actual OID type.

[0-9]:.*cannot be imported from module—MIB failed to import because order is not being defined properly.

Routing Protocols

- **Option to display routing instance table in the show route advertising-protocol output**—Beginning with Junos OS Release 16.2, you can use the **show route advertising-protocol table foo** command to display the routing instance table for any address family on a VPN route reflector, or a VPN AS boundary router that is advertising local VPN routes . However, if you do not specify the table in the command, the output displays each VRF prefix twice.
- **Timers of delay-route-advertisements are modified**—Beginning with Junos OS Release 15.1F7, the range of the timer values of **delay-route-advertisements** has been increased to 36000 from 3600. The default value of **route age**, that is the maximum delay after route aggregates have been created has also been modified to 0. In earlier Junos releases, the default **route age** was 1200. The timer values of **delay-route-advertisements** are configured to avoid premature route advertisements that might result in traffic loss in a BGP session.

  [See delay-route-advertisements.]
- **Change in default behavior of router capability (MX Series)**—Starting in Junos OS Release 15.1F7, 16.1R4, 16.2R2, 16.1X65, and 17.1R1 and later , router capability TLV distribution flag (S-bit), which controls IS-IS advertisements, will be reset, so that the segment routing capable sub-TLV is propagated throughout the IS-IS level and not advertised across IS-IS level boundaries.

Services Applications

- **Change in option name to configure inactive timeout for IKE ALG child sessions (MX Series)**—Starting in Junos OS Release 16.2R1, the name of the option to configure the inactive timeout for Internet Key Exchange (IKE) application-level gateway (ALG) child sessions is changed from **child-session-timeout** to **child-inactivity-timeout**.
- **Deprecated security idp statements (MX Series)**—Starting in Junos OS Release 16.2R2 and earlier, the **[edit security idp]** configuration statements are deprecated for the MX Series.
- **Change in behavior of IKE negotiation (MX Series)**—Starting in Junos OS Release 16.2R2, when you commit an IPsec configuration that includes **establish-tunnels immediately** at the **[edit services ipsec-vpn]** hierarchy level, the service set might take up to 30 seconds to initiate IKE negotiations.

Subscriber Management and Services

- **Configuring a pseudowire subscriber interface for a logical tunnel (MX Series)**—Starting in Junos OS release 16.1R2, 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.

- **Change in range for PPP keepalive interval (MX Series)**—Starting in Junos OS Release 16.2, 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 nonsubscriber services are handled by the Routing Engine with an interval range from 1 second through 32,767 seconds.

- **Automatic limit set for transmit window size (MX Series)**—Starting in Junos OS Release 16.2, 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.

- **New option to display all pending accounting stops (MX Series)**—Starting in Junos OS Release 16.2R2, the **brief** option is added to the **show accounting pending-accounting-stops** command. This option displays the current count of pending RADIUS Acct-Stop messages for subscribers, services, and total combined value:

  ```
  user@host> show accounting pending-accounting-stops brief
  ```

  ```
  Total pending accounting stops: 4
    Subscriber pending accounting stops: 2
    Service pending accounting stops: 2
  ```

- **DNS servers displayed by the show subscribers extensive command (MX Series)**—Starting in Junos OS Release 16.2R2, the output display of DHCP domain name servers (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, IPv6 Secondary DNS Address.

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

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

- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 16.2R2, the **show subscribers extensive** command displays the IPv6 Interface Address field only when the dynamic profile includes the $junos-ipv6-address predefined variable.

  In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

- **Change to DHCP option 82 suboptions support to differentiate duplicate clients (MX Series)**—Starting in Junos OS Release 16.2R2, only the ACI (suboption 1) and ARI (suboption 2) values from the option 82 information are considered when this information is used to identify unique clients in a subnet. Other suboptions, such as Vendor-Specific (suboption 9), are ignored.

- **Default L2TP resynchronization method changed and statement deprecated (MX Series)**—Starting in Junos OS Release 16.2R2, 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 that it is deprecated.

- **Traffic shaping and L2TP tunnel switches (MX Series)**—Starting in Junos OS Release 16.2R2, 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 L2TP network server (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.

- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 16.2R2, the **show subscribers extensive** command displays the IPv6 Interface Address field only when the dynamic profile includes the $junos-ipv6-address predefined variable.

  In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 16.2R2, the **show subscribers extensive** command displays the IPv6 Interface Address field only when the dynamic profile includes the $junos-ipv6-address predefined variable.

  In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

- **IPv6 Link Local Addresses Assigned to Underlying Static Demux Interfaces (MX Series)**—Starting in Junos OS Release 16.2R2, 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
      }
    }
  ```

- **Change in display of IPv6 Interface Address field by the show subscribers extensive command (MX Series)**—Starting in Junos OS 16.2R2, the **show subscribers extensive** command displays the IPv6 Interface Address field only when the dynamic profile includes the $junos-ipv6-address predefined variable.

  In earlier releases, the command always displays this field, even when the variable is not in the profile. In this case, the field shows the value of the first address from the Framed-IPv6-Prefix attribute (97).

- **Memory mapping statement removed for Enhanced Subscriber Management (MX Series)**— In Junos OS Release 17.3R1, use the following command when configuring database memory for Enhanced Subscriber Management:

  set system configuration-database max-db-size

  CLI support for the **set configuration-database virtual-memory-mapping process-set subscriber-management** command has been removed to avoid confusion. Using the command for subscriber management now results in the following error message:

  WARNING: system configuration-database virtual-memory-mapping not supported.
  error: configuration check-out failed.

  [See Interface Configuring Junos OS Enhanced Subscriber Management for an example of how to use the **max-db-size** command.]

User Interface and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (MX Series and T Series)**—Starting in Junos OS Release 16.2R2, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not

enclosed in quotation marks. Prior to Junos OS Release 16.2R2, integers in JSON configuration data were treated as strings and enclosed in quotation marks.

- **Changes to the XML and JSON output when displaying the differences between the candidate and active configurations (MX Series)**—Starting in Junos OS Release 16.2R2, when you compare the candidate and active configurations and display the differences in XML or JSON format, for example by using the **show | compare | display (json | xml)** CLI command or the **<get-configuration compare="rollback" format="(json | xml)">** RPC, the device omits the **<configuration>** tag in the XML output and omits the **configuration** object in the JSON output if the comparison either returns no differences or if the comparison returns differences for only non-native configuration data, for example, configuration data associated with an OpenConfig data model.

VLAN Infrastructure

- **ACI and ARI from PADI messages included in Access-Request messages for VLAN authentication (MX Series)**—Starting in Junos OS Release 16.2, 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

- **Support for ping on a virtual gateway address**—Starting in Junos OS Release 16.2R2, Junos supports pinging an IPv4 or IPv6 address on the preferred virtual gateway interface. To set up support for ping, you must include both the **virtual-gateway-accept-data** and the **preferred** statements at the **[edit interfaces irb unit]** hierarchy of the preferred virtual gateway. This enables the interface on the preferred virtual gateway to accept all packets for the virtual IP address, including ping packets.

See Also

- New and Changed Features on page 25
- Known Behavior on page 52
- Known Issues on page 54
- Resolved Issues on page 62
- Documentation Updates on page 115
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 124

Known Behavior

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 16.2R2 for 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.

High Availability (HA) and Resiliency

- **ISSU limitations with MX Series FPCs**—Unified ISSU is not supported with Junos OS Release 16.2 for MX Series routers with the following Flexible Port Concentrators (FPCs):
  - MX-MPC2E-3D-NG
  - MX-MPC2E-3D-NG-Q
  - MX-MPC3E-3D-NG
  - MX-MPC3E-3D-NG-Q

  To perform a unified ISSU on a MX Series router with these FPCs installed, reboot the FPCs before completing the unified ISSU process.

  [See Unified ISSU System Requirements]

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.

Software Installation and Upgrade

- On a broadband network gateway (BNG) that is running enhanced subscriber management, you must take the service cards offline before you can perform an in-service software upgrade (ISSU) to Junos OS Release 16.2 from a Junos OS release that includes the application-aware policy control feature (Junos OS Release 16.1R4 and later).

Subscriber Management and Services

- For dual-stacked clients over the same PPP over L2TP LNS session, enhanced subscriber management does not support configurations where both of the following are true:
  - The CPE sends separate DHCPv6 solicit messages for the IA_NA and the IA_PD.
  - The solicit messages specify a type 2 or type 3 DUID (link-layer address).

As a workaround, you must configure the CPE to send a single solicit message for both IA_NA and IA_PD when the other configuration elements are present.

See Also

- New and Changed Features on page 25
- Changes in Behavior and Syntax on page 43
- Known Issues on page 54
- Resolved Issues on page 62
- Documentation Updates on page 115
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 124

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 16.2R2 for 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 54
- General Routing on page 56
- High Availability (HA) and Resiliency on page 59
- Interfaces and Chassis on page 59
- Junos Fusion Provider Edge on page 59
- Layer 2 Features on page 59
- MPLS on page 60
- Platform and Infrastructure on page 60
- Routing Protocols on page 61
- Services Applications on page 61
- Subscriber Access Management on page 61
- VPNs on page 62

Forwarding and Sampling

- Impact of Policing filter creation and application to the LSP in the same commit sequence—It is known that policing filter application to the LSP is catastrophic. Any active LSP carrying traffic when applied a policing filter tears down and resignals and drops traffic for ~2 seconds. In Junos OS Release 16.1R1, it would take up to 30 seconds for the LSP to come up if 1. Creation of the policing filter and application of the same to the LSP through configuration in the same commit sequence 2. Load override of a configuration file that has policing filter and policing filter application to the LSP followed by commit. PR1160669

- Inline JFlow (MXVC) : NextHop Address/OIF being reported by IPv6 template on MXVC setup is incorrect–Root Cause of the Problem: +++++++++++++++++++++++++ As per the investigation from RPD : we have is an interface for a direct route starting in ifdown condition. The remote side is then brought up, so I/F goes to ifup. Since it is a direct route, rpd does not install the route or nexthop. It receives that info from the kernel, and just updates a nexthop in rpd local storage. route and nexthop for the interface are taken care of in the kernel. There is no route change in rpd. route_record depends on route flash to find out about updates. That is the architecture. Since there is no route change, there is no route flash, so route_record is blissfully unaware. In order to change this, we would need to decide that we want a route flash for this case. Currently, for direct and local routes / nexthops, these are "don't care" in rpd, as far as route updates go. We just update our nexthop info, without marking for any other notifications. To change this, we would need to find the correct place to decide we need to flash the route, and at the same time, make sure we don't do any harm to anything else. That is what I am currently working on finding. A complication for the solution is a change that was done for PR 1002287, where if the NOTINSTALL flag is set, do not send the update to srrd. That flag is set for direct and local routes. Incidently, this is day-one operation. If the interface is up at startup, it should all work correctly. Why is the pfe depending on rpd / srrd to get the info for sampling when it is already there in the forwarding table ? +++ ++ FIB table can provide OIF/GW only. SRC_MASK, DST_MASK, SRC_AS and DST_AS are not available in PFE FIB Table. So SRRD connection is required. Listening to both SRRD and FIB table, and consolidating information will complicate implementation. Scanning entire FIB Table just for the few such routes will have performance impact and will complicate present implementation. This is day 1 implementation for SRRD/Sampled. Workarounds: ++++++++++++ There are two possible workarounds a) A workaround would be to have the far end interface up when the DUT interface is brought up. In the case where that is not happening, a recovery would be to disable the DUT interface, then enable it again. At that point, everything should be initially brought up in the state we are looking for. b) enable nexthop-learning knob. Please refer to the documentation on the working of this knob before enabling. PR1224105

- The "dfwinfo: tvptest:dfwlib_owner_create tvp driven policer_byte_count support 0" message is seen after running the **show firewall** command. This behavior is specific to Junos OS release 16.1 and is a cosmetic issue. << sample config >> set interfaces ge-0/0/0 unit 0 family inet filter input test_filter set interfaces ge-0/0/0 unit 0 family inet address 100.100.100.1/24 set firewall family inet filter test_filter term policer then policer policer_test set firewall policer policer_test if-exceeding bandwidth-limit 100m set firewall policer policer_test if-exceeding burst-size-limit 125k set firewall policer policer_test then loss-priority low PR1248134

General Routing

- The static subscriber process (jsscd) might crash in a scaled environment -- The jsscd process might crash in a 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 1 11:11 /var/tmp/jsscd.core-tarball.0.tgz. PR1133780

- On MX routers simultaneously equipped with DPC and MPC, the following messages are reported by MPC and DPC respectively, and traffic loss might be observed after performing ISSU.

  MPC reports "FI Cell underflow at the state stage" DPC reports "Non first cell drops in ichip fi rord:xxxx"

  The issues is seen on the following MPCs:

  - MX-MPC1-3D
  - MX-MPC1-3D-Q
  - MX-MPC1E-3D
  - MX-MPC1E-3D-Q
  - MX-MPC2-3D
  - MX-MPC2-3D-Q
  - MX-MPC2-3D-EQ
  - MX-MPC2E-3D
  - MX-MPC2E-3D-Q
  - MX-MPC2E-3D-EQ
  - MPC-3D-16XGE-SFPP
  - MPCE-3D-16XGE-SFPP

  PR1163776

- Chef for Junos supports additional resources to enable easier configuration of networking devices. These are available in the form of netdev-resources. The netdev-resource developed for interface configuration has a limitation to configure XE interface. Netdev-interface resource assumes that 'speed' is a configurable parameter which is supported on a GE interface but not on an XE interface. Hence netdev-interface resource cannot be used to configure an XE interface due to this limitation. This limitation is applicable to packages chef-11.10.4_1.1.*.tgz chef-11.10.4_2.0_*.tgz in all platforms {i386/x86-32/powerpc}. PR1181475

- AMS redundant interfaces not listed under possible- auto-completions the following operational commands:

  - **show interfaces redundancy**
  - **request interface switchover**

PR1185710

- AMS redundant interfaces not listed under possible auto-completions for the following operational commands:
  - **show interfaces redundancy**
  - **request interface switchover**

  PR1185710
- 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 the up state. PR1192161
- During ISSU (unified in-service software upgrade) it might be noticed that the below log messages are displayed: SFP: pointer Null, sfp_set_present This might trigger a flap in the interfaces on MX routers while upgrading using ISSU. PR1200045
- A few sessions are always dropped during session setup with IPsec, consistently seen above 1M sessions. PR1204566
- Major errors might be seen on MPC3/FPC3 with 1X100 and 5x100 DWDM MIC/PIC. user@router> show chassis alarms no-forwarding 1 alarms currently active Alarm time Class Description <timestamp> Major FPC 3 Major Errors The following messages are seen in the logs: fpc3 Cmerror Op Sub Set: 5-port 100G DWDM MIC/PIC : 5-port 100G DWDM MIC/PIC(3/0) link 0 : DSP loss of lock fpc3 Cmerror Op Sub Set: 5-port 100G DWDM MIC/PIC : 5-port 100G DWDM MIC/PIC(3/0) link 0 : DFE tuning failed alarmd[16241]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 3 Major Errors craftd[15906]: Major alarm set, FPC 3 Major Errors. PR1204566
- The /etc/passwd file is created in the process of the first commit when a pristine jinstall image is used to boot for the first time. If event-options is configured, the system will try to read the configuration from the available event scripts which requires privileges obtained from the /etc/passwd file. That causes a circular dependency as the commit will not pass if the configuration includes event-options the first time a pristine image boots up, which is the case of an upgrade performed with virsh create. PR1220671
- No optic lane diag exported for XFP optic in both CLI and snmp. PR1223742
- A wrong PE is being attached to an ESI when the router receives two copies of the same AD/ESI route (e.g. one through eBGP and another one received from an iBGP neighbor). This will causes partial traffic blackhaule 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: 00:13:78:00:00:00:00:00:00:01 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
- When SW detects an uncorrectable XR2 error, which are in fixed locations relative to queues in XQ, it removes the queues from service by moving the traffic to a new L4 and new set of queues, using other XR2 locations. Currently the queues/L4 that are removed are never returned to service until reset. This mechanism would be expanded to include L4NP parity errors, and possibly others as well. In this case when an L4NP parity error is detected we remove the L4 and queues from service. PR1232952

- On MX series with rpd in "ASYNC" mode, if the distributed IGMP is configured, rpd core might be seen, and causing rpd crash. PR1238333
- For ANCP subscribers in Idle state the previously reported speed in ANCP Port UP message is not applied. PR1242992
- ANCP neighbors going down after commit in case any ANCP related configuration was changed. PR1243164
- In a junos telemetry interface environment, the FPC might crash when adding physical interface sensor during FPC coming up. It is because during FPC coming up, the IFD (physical interface) does not exist. At this point, if accessing IFD, the FPC might crash. PR1243411
- VPLS mac table is not being populated properly when checked with CLI "show vpls mac-table", though all subscribers have traffic. Thus it is considered a cosmetic issue. PR1257605
- On some T series platform routers, the LSI statistics are not shown in Aggregated Ethernet Interface Bundles and also the input stats counter for Aggregated Ethernet does not include MPLS traffic. PR1258003
- Due to transient Hardware error conditions only syslog events XMCHIP(x) FI: Cell underflow at the state stage - Stream 0, Count 65535 are reported which is a sign of fabric stream wedge. Additional traffic flow register pointers are validated and if stalled a new CMERROR alarm is raised "XMCHIP(x) FI: Cell underflow errors with reorder engine pointers stalled - Stream 0, late_cell_value 65535, max_rdr_ptr 0x6a9, reorder_ptr 0x2ae" PR1264656
- Due to transient Hardware events, fabric stream may report 'CPQ1: Queue unrderrun indication - Queue <q#>' in continuous occurrence. For each such events, all fabric traffic is queued for this Packet Forwarding Engine reporting the error and causes very high amount of fabric drops.PR1265385
- **HALP-lbnh_xlate_cntr_db_get_stats:250counter id 1573873: Unable to find lbnh xlate counter** messages are flooding the syslog. This is only a syslog entry and there is no operational impact. PR1268452
- On MIC-3D-20GE-SFP-E and MIC-3D-20GE-SFP-EH, an interrupt threshold was introduced, if MIC error interrupts are more than the threshold (> 2500 per 5min), the MIC will be restarted. Due to that change, MIC error interrupts will hog the CPU when restart is initiated. PR1270420

High Availability (HA) and Resiliency

- In a rare scenario, GRES might not reach the ready state and might fail to start, because the Routing Engine does not receive the state ack message from the Packet Forwarding Engine after performing graceful Routing Engine switchover (GRES). This is a timing issue. It might also stop Routing Engine resource releasing and then cause resource exhausting. Reboot the system if this problem occurs. PR1236882

Interfaces and Chassis

- After changing the MTU on the IFD, on the static vlan demux interface above the IFD IPv6 Link Local address is not assigned. PR1063404
- During configuration changes and reuse of Virtual IP on an interface as a interface address, it is required to delete the configuration do a commit and then add the interface address configuration in the following commit. PR1191371
- IPV6 neighbor ship is not created on the IRB interface. PR1198482
- In a VPLS Multi-homing scenario, the CFM packets are forwarded over the standby PE link resulting in duplicate packets or loop between the active and standby link. PR1253542
- JUNOS upgrade involving releases 14.2R5 (and above in 14.2 maintenance releases) and 16.1 above mainline releases with CFM configuration can cause CFMD core post upgrade. This is due the old version of /var/db/cfm.db. PR1281073

Junos Fusion Provider Edge

- In a Junos Fusion setup, fixed connectivity issues between two VLANs on the same extended port. PR1264900
- Junos OS upgrade involving Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration can cause cfmd to generate a core file after upgrade. This is due to the old version of **/var/db/cfm.db**. PR1281073

Layer 2 Features

- When **input-vlan-map** with a push operation is enabled for dual-tagged interfaces in "Enhanced-IP" mode, there is a probability that the broadcast, unknown unicast, and multicast (BUM) traffic might be silently dropped or discarded on some of the child interfaces of the egress Aggregated Ethernet (AE) interfaces or on some of the equal-cost multipath (ECMP) corelinks. PR1078617
- Starting in Junos OS Release 14.2R3 the **show class-of-service fabric statistics** CLI command might fail with **Error = Operation timed out** message in some cases (especially if there are many FPCs in the chassis). This occurs because data structures used to query fabric statistics became significantly larger in later releases. Thus when multiple FPCs start transmitting data to the Routing Engine at the same time, some packets might get dropped in the internal Ethernet switch on the control board. If retransmission does not happen within the timeout, the **Operation timed out** error is seen. PR1228293

- After changing the underlying physical interface (IFD) for a static VLAN demux interface the NAS-Port-ID is formed still based on the previous IFD. PR1255377
- In VPLS topologies the kernel might report the error **pointchange for TLV type 00000052 not supported on IFL <name>** in **/var/log/messages** where <name> is a VT or LSI interface used by VPLS. The trigger to cause the issue depends on timing and is most often seen with high VPLS pseudowrite scaling when multihoming is configured, but other triggers might apply as well. The problem might cause high rpd CPU utilization, which can slow routing convergence. PR1279192

MPLS

- 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 will need a manual restart with **restart routing**. PR1238698
- A new configuration **protocols mpls traffic-engineering bgp-igp-both-ribs** in the routing instance is required to make cOC work. PR1252043
- The throughput measurement might be inaccurate when doing performance measurement on an MPLS label-switched path. PR1274822
- The throughput measurement may be inaccurate when doing performance measurement on a MPLS label-switched-path. PR1274822

Platform and Infrastructure

- When TCP authentication is enabled on a TCP session, the TCP session may not use the selective acknowledgement (SACK) TCP extensions. PR1024798
- On MX Series platform, parity memory errors might happen in pre-classifier engines within a MPC. Packets will be silently discarded as such errors are not reported and makes it harder to diagnose. After the change in this PR, CM-ERRORs, such as syslogs and alarms, will be raised when parity memory errors occur. PR1059137
- SNMP queries to retrieve jnxRpmResSumPercentLost will return the RPM/TWAMP probe loss percentage as an integer value, whereas the precise value (including decimal points) can be retrieved through the CLI by using the following commands: **show services rpm probe-results**, and **show services rpm twamp client probe-results**. PR1104897
- The mustd daemon might crash when large configurations are committed. PR1186326
- Multicast traffic might get dropped when the STP port role is changed. As a workaround, toggle the IGMP snooping membership. PR1193325
- On rare occasions during the route add/delete/change operation, the kernel might encounter a crash with the error **rn_clone_unwire no ifclone parent**. PR1253362

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 the context of large number of configured VPNs, routes changing in the midst of a bgp path-selection configuration change can sometimes lead to an rpd core. This core has been seen with the removal of the **always-compare-med** option. PR1213131
- Starting in Junos OS Release, 16.1R2, when BGP add-path is configured and the same prefix is received from multiple peers with different source AS, depending on the order that the prefix advertisements are received in, the rpd might crash. PR1223651
- On rpd crash with switchover-on-routing-crash enabled on box, live vmcores might be seen on both Routing Engines without an impact on the system. PR1267796

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 keep too busy to bring up new sessions. PR990081
- When MS-PIC is running on T640/T1600/T4000, the number of maximum service sets is incorrectly limited to 4000, instead of 12000. This might impact a scaled service (IPsec, IDS, NAT, stateful firewall filter, and so on) environment. PR1195088
- When loading or rolling back a configuration that removes a serviceset and changes where the MS interfaces are assigned, traffic might be silently dropped or discarded to a series of the existing service sets. PR1223302
- If an L2TP subscriber has static pp0 interface on the LAC side, LCP renegotiation is configured on the LNS side and the CPE has been changed, it can cause an issue with successful negotiation of the PPP session between LNS and CPE. PR1235554
- Account Session ID, Interface Identifier, and Subscriber User Name trigger attributes are optimized for a scaled subscriber management environment. If you include any of the other, non-optimized, trigger attributes in a scaled subscriber management environment, a significant delay might be observed between the time when the DTCP ADD message is sent and the time when forwarding starts for the mirrored traffic. For example, if there are 10,000 subscriber sessions on the router, forwarding of the mirrored traffic might be delayed for 20 minutes. This delay occurs when you specify any non-optimized attribute, with or without any optimized attribute. The delay occurs regardless of the order of attributes in the DTCP packet. PR1269770

Subscriber Access Management

- On MX series platform, when using RADIUS dynamic requests for subscriber access management, if the device detects that the CoA-Request it received is he same as one already being processed, the device incorrectly sends the CoA-NAK packet back to the

RADIUS server with incorrect code 122 (invalid request), before sending the 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

- 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
- Subscribers get stuck in terminated state during pppoe login/logout test. PR1262219

VPNs

- In NG-MVPN scenario, when "forwarding-cache timeout never non-discard-entry-only" is configured for an MVPN instance, even though the cache lifetime is shown as forever in the output of CLI command "show multicast route instance X extensive", the route disappears after 7-8 minutes. PR1212061

See Also

- New and Changed Features on page 25
- Changes in Behavior and Syntax on page 43
- Known Behavior on page 52
- Resolved Issues on page 62
- Documentation Updates on page 115
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 124

Resolved Issues

This section lists the issues fixed in the Junos OS main 16.2R2 Release for 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.

- Resolved Issues: 16.2R2 on page 63
- Resolved Issues: 16.2R1 on page 99

Resolved Issues: 16.2R2

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 could be seen by executing the CLI command "show system core-dumps". PR1066009
- In phase 1 of Junos Fusion Provider Edge, extended ports do not support SNMP-based CoS statistics. Polling an EP port for CoS stats can trigger a cosd core file. PR1205512
- When CoS is configured, in a very rare situation, due to the timing issue between dcd and cosd during commit, the cosd might crash. For example, if you delete an interface that belongs to an AE interface and then configure it as a single port with CoS in a single commit, this issue might occur. PR1220524
- The "show interfaces queue <if-name>" command has three display options:
  - show interfaces queue <if-name>: Displays queued/transmitted/dropped packets/bytes for all IFD children.
  - show interfaces queue <if-name> aggregate: Displays queued/transmitted/dropped packets/bytes for all IFD children except for IFD RTP traffic
  - 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.

  The following should be true for queues depth values:

  - show interfaces queue <if-name>: Displays queues depth values for RTP queues
  - show interfaces queue <if-name> aggregate: Displays queues depth values for RTP queues
  - 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 MX Series and T Series routers with ingress and egress queueing enabled, input traffic-control-profile is configured, but no output traffic-control-profile on IFL. After you activate/deactivate the CoS configuration, the cosd process might crash. PR1236866
- The error message of cos_check_temporal_buffer_status might be observed when configuring hierarchical CoS with strict-high scheduling. PR1238719

Forwarding and Sampling

- On MX Series routers, if the "compress" and "backup-on-failure" options are configured under accounting-options file <file-name> hierarchy, due to an issue in the source file lookup when "compress" option is enabled, local backup might not perform after transfer to archive site fails. PR1198095
- If a two-color policer is configured on MX Series routers with MPCs/MICs, more traffic than the limited traffic might be passed when packet 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
- In firewall_service.proto file, AccessListObjBind changed. The member "bind_object" in AccessListObjBind is no longer a string object; it changed to "one of" structure which is shown as follows:

  message AccessListObjBind { // ACL AccessList acl = 1; // Binding object type AccessListBindObjType obj_type = 2; // Bind object name where the ACL is to be bound string bind_object = 3; // Bind direction AclBindDirection bind_direction = 4; // Family on the bind object. Must match with the ACL family AccessListFamilies bind_family = 5; }

  Starting in Junos OS Release 16.2R2 release AccessListObjBind message member “string bind_object“ changed as follows:

  message AccessListBindObjPoint { oneof OneOf_AclBindPoint { // Bind object name where the ACL is to be bound string intf = 1; } } /* * Per forwarding element ACL binding */ message AccessListObjBind { // ACL AccessList acl = 1; // Binding object type AccessListBindObjType obj_type = 2; // Bind object name where the ACL is to be bound - string bind_object = 3; + AccessListBindObjPoint bind_object = 3; // Bind direction AclBindDirection bind_direction = 4; // Family on the bind object. Must match with the ACL family AccessListFamilies bind_family = 5; }PR1230587
- When a firewall filter (family "any") with a shared-bandwidth-policer is applied on an MC-AE interface, it will be configured with bandwidth 0 and carve-up factor 0 as expected. But after MC-AE A/S switchover when standby becomes active, the policer would not reconfigure, still have the bandwidth of 0 and drop all packets. PR1232607
- With sampling configuration, if you do not define a version for the second flow server, after committing configuration, the backup Routing Engine might reboot. It might affect how routing protocols are replicated to the backup Routing Engine. PR1233155
- On MX Series routers with "ipv4-flow-table-size" or "ipv6-flow-table-size" configuration, if sampling instance is not defined under chassis hierarchy (sampling instance is not associated to FPC), after rebooting the router, the "ipv4-flow-table-size" or "ipv6-flow-table-size" does not propagate to FPC. PR1234905
- When 'push-backup-to-master' knob is configured under accounting-options file section, the corresponding accounting files need to be pushed to master RE from standby RE. But due to a software defect, the following issues are observed.
  - 1) The files push from standby Routing Engine to master Routing Engine was happening irrespective of this **push-backup-to-master** configuration statement.

- 2) The files push from standby Routing Engine to master Routing Engine was not happening when the backup option is configured as 'master-only'.

  PR1236618

- J-Flow version 9 cannot get TCP flag information from IPv6 fragment packets. However, it can get other information like src and dst ports infromation. It can get sampling information partially from the TCP header in IPv6 fragment packets. PR1239817
- J-Flow version 9 is sending the flows with the source-address inverted in the show firewall log. PR1249553
- On MX Series routers, after GRES or configuration change that leads to pfed core file and restart, the routers might send for every single session 5 AcctInterim update. PR1249770
- In MX Series subscriber management environment; the layer 2 address learning daemon (l2ald) daemon might crash during EVPL subscriber login logout stress test. PR1258853
- The final service stats are queried via the on-demand service stats handling module of the pfed process. When the responses are returned from the Packet Forwarding Engine to the Routing Engine through pfed), they are mapped to the request via the request ID as well as location offset. When there are more than one filter configured for a BBE filter service session (out of IPV4,IPV6 IN,OUT filters), more than one request will be sent to the same location (Packet Forwarding Engine) with the same request ID. PR1262876
- Routing-instances information of the physical interface is not showing in the flat accounting file when the interface is attached to the aggregate Ethernet interface. This behavior is seen when using flat file accounting for L2BSA subscribers. PR1275225

General Routing

- This is a timing issue. After deleting and reconfiguring a VRF instance or changing route-distinguisher in VRF instance while rpf-check is enabled, the rpd process might crash. The routing protocols are impacted and traffic disruption will be seen due to loss of routing information. PR911547
- In an MX Series Virtual Chassis (MX-VC) environment, the private local next hops and routes pointing to private local next hops are sent to the Packet Forwarding Engine from the master Routing Engine and not to the secondary Routing Engine. Next, a Routing Engine switchover happens. Because the new master Routing Engine does not detect such next hops and routes, they are not cleaned up. When a next hop with the 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. PR951420
- An incorrect byte count was seen in the ipfix exported statistics packets for MPLS flows. PR1067084
- When ingress and egress layer2-overhead is configured on a dynamic subscriber interface, the layer2-overhead bytes are not added to the IFL stats. PR1096323
- On MX Series routers with MS-MPC/MS-MIC in use, if the NAT session is freed/removed without removing the timer wheel entry, the MS-MPC/MS-MIC might crash. This is a

timing issue in which just before invoking the timer wheel callback, the NAT session extension got freed/removed. PR1117662

- With l2tp subscribers, all FPCs except the card that hosts subscribers will report a log message "jnh_if_get_input_feature_list(9723): Could not find ifl state" after every subscriber's login attempt. PR1140527
- On MX Series routers with services PIC (MS-DPC/MS-MPC/MS-MIC), the ICMP time exceeded error packet is not generated on an IPsec router on the de-encapsulation side. PR1163472
- On MS-MPC or MS-PIC, OSPF adjacency may fail to establish when there is no static route pointing to service PIC. PR1164517
- In all Junos OS with EVPN scenario, the Layer 2 address learning daemon (l2ald) might not clean up the RNH_LE entry when the BGP neighbor is down and cause the end-to-end traffic of EVPN to be dropped. PR1173420
- NAT64 service-set:Port block efficiency and unique pool users statistics display incorrect values when the NAT POOL is modified dynamically with CGNAT traffic for the particular term in the NAT rule. PR1177244
- On dual Routing Engine systems, the false alarm message "Host 1 failed to mount /var off HDD, emergency /var created" is observed if the master Routing Engine and backup Routing Engine are running on different Junos OS versions. PR1177571
- 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 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 an MX Series router, PIC card on MS-MIC/MS-MPC might crash in rare cases. This is a timing issue that might cause traffic loss and has no exact aspect of configurations for trigging the issue. PR1182807
- Fragmented ALG control traffic is not supported on the MS-MPC or MS-MIC. PR1182910
- On MX Series routers, MS-MIC crash might occur. The exact trigger of the issue is unknown; generally, this issue might happen very rarely without any external triggers. The crash might occur with any services configuration, with core files pointing to a Program terminated with signal 4, Illegal instruction. PR1183828
- FRU model numbers might be missing or incorrect as follows.

  740-013110 PDM-MX960

  740-057995 FFANTRAY-MX960-HC-S

  750-033205 MX-MPC3E-3D (incorrect)

  750-038493 MX-MPC2E-3D-Q

  750-044130 MX2K-MPC6E

  750-045372 MX-MPC3E-3D

750-046005 MPC5EQ-100G10G

750-046532 MIC6-10G

750-049457 MIC6-100G-CFP2

750-054563 MPC5E-40G10G

750-054902 MPC3E-3D-NG

750-054903 MPC2E-3D-NG-Q

750-055976 SCBE2-MX-S

all CFP, CFP2, QSFPP, QSFP28 optics

all MX2000 FRUs

all MPC7E, MPC8E, MPC9E, SFB2 FRUs

Note that 'show chassis hardware models' displays correct information, but optics are missing from that output. PR1186245

- On a Junos-based platform, CHASSISD_I2CS_READBACK_ERROR error might occur on a single occurrence of I2C read failure. A single occurrence is a transient error and may be seen randomly without any particular trigger. This type of message should be reported 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
- On MX Series routers with NAT service configured on AMS interfaces, after rebooting FPC/PIC, the NAT pool split between AMS members is incorrect. There are overlapping IP pools and sometimes missing pools, causing NAT to work incorrectly. PR1190461
- On MX Series routers with Junos Telemetry Interface, and with the "set routing-options lsp-telemetry" configuration statement configured. When SDN-telemetry (the agentd process) is disabled or continuously restarted, certain messages are repeatedly logged into syslog, the rpd and eventd processes CPU may get near 100%, and eventually the agentd also gets near 100%. When this issue happens, the agentd process is not able to accept new subscriptions, dropping all existing subscriptions. It can be triggered by restarting consecutively SDN-telemetry (the agentd process), or after device reboot. PR1192366
- In an MX BNG subscriber management environment, Radius accounting statistics provided by the MX Series BNG might slightly deviate from the actual statistics if the subscriber session terminated abruptly while traffic flow was active. PR1192775
- Configuring an RLT interface and rebooting the router shows the RLT interface is down. The show l2circuit connection shows an MTU mismatch as the immediate cause. PR1192932
- Prior to this PR, when T-series SCG lost an external clock source, clock state remained hold-over mode forever. This PR has changed the behavior so that the state would automatically be changed from hold-over to free-run after 24 hours. PR1197380

- On MX Series routers with MPC5E installed, in a high-temperature situation, the temperature thresholds for triggering the high temperature alarm and controlling fan speed are based on the FPC level. Any sensor values in the FPC that exceed the temperature threshold of the FPC trigger the actions associated with temperature thresholds. PR1199447
- With MPC8/9 MRATE MIC. With a plug-in optics module(QSFP28-100GBASE-LR4), bit errors might be seen. PR1200010
- On MX Series routers, the mspmand process might crash on the MS-MPC with XLP B2 chip (for example, REV17). The exact trigger is unknown. It is usually seen with 70% to 90+% CPU load conditions. PR1200149
- When performing unified ISSU on MX Series routers, the MPC might crash during the field-replaceable unit (FRU) upgrade process. PR1200690
- A dynamic tunnel gets timed out every 15 mins by default, and then re-tries to create another tunnel. This happens if the route obtained from IGP is non-forwarding. PR1202926
- When PPPoE subscribers log in to or out of the device, an SNMP link up/down trap will be generated by the system if "no-trap" is configured in the corresponding dynamic-profile. PR1204949
- SMID daemon has stopped responding to the management requests after a jl2tpd (L2TP daemon) crash on a production MX960 BNG. PR1205546
- Problem - 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 continous IIF_mismatches. PR1206121
- In an L2TP scenario, in a rare situation, the command "show subscribers summary port extensive" output might have an incorrect tunneled/terminated sessions count due to an issue with populating the outputs. There is no traffic impact. PR1206208
- When PCEP is enabled and LSPs are undergoing changes, like make before break (MBB) for rerouting, the rpd has to send those updates to the PCE. However, when the PCEP session to PCE goes down, these updates are cancelled, but the rpd fails to completely reclaim the memory allocated for these updates. This causes increases 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 to display chassis components, the Compact Flash card and hard disk serial numbers may be truncated to 15 characters. PR1209181
- On MX Series routers, if any inline feature is configured (for example, inline BFD, CFM , and PPP), the FPC might crash and core files are generated. PR1210060

- The Periodic Packet Manager (ppman) based sessions (such as CFM session) might be flapping when executing offline/online MIC-3D-20GE-SFP (model number) MIC inserted into MPC2E-NG/MPC3E-NG. This occurs because the TNPC-CM thread is hogging the CPU for ~450 ms when executing MIC-3D-20GE-SFP MIC offline/online. PR1211702
- When an ARP entry is learned through the Aggregated Ethernet interface, and a route is pointing to that ARP next hop, the ARP entry might not expire even though the ARP IP is no longer reachable. This issue is due to the route next hop on the AE interface getting stuck in unicast state even if the remote end is not reachable, and the RPD never gets to determine that ARP is invalid. The route nexthop on Aggregated Ethernet interface should be shown in 'hold' state when the remote end is not reachable. PR1211757
- On EVPN/VXLAN setup with the MX Series router as PE device, when both arp aging-timer and static MAC applied on the IRB interface associated with EVPN, the packet originating from Routing Engine on the PE router (such as ping) to the core side might be corrupted. This issue only impacts the traffic originated from the Routing Engine and does not impact the transit traffic. PR1213062
- On MX Series routers with MPC3/MPC4/MPC5/MPC6/MPC2-NG/MPC3-NG line cards, the chassisd process crashes continuously on both Routing Engines because some failure cases caused by underlying software and hardware are not handled gracefully. Both Routing Engines might lost mastership and get stuck in backup mode. PR1213808
- If a 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
- 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
- In large-scale configurations or environments with high rates of churn, MX Series routers with FPC's ASIC memory will become "fragmented" over time. In an extreme case, it is possible that memory of a particular size will become exhausted. Also, due to the fragmentation, the available memory will not fulfill the pending allocation. PR1216300
- When VPLS instances are configured for the first time or when a system with VPLS instances is rebooted, rpd will be consuming high CPU usage (100%) for a period (10-20 mins), the installation of other routes may defer and traffic will be lost. Many other RPD services may also slow down or be unavailable. PR1216332
- 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. The logs are harmless and can be ignored. PR1217975
- On MX Series routers, replacing an MQ FPC (MPC Type1, 2, MPC 3D 16x10GE) with an XM one (MPC Type 3,4,5 6. 2E-NG, 3E-NG) might cause all other MQ-based cards to report "FI Cell underflow at the state stage". It will cause packets to be dropped. PR1219444
- If RS/RA messages were received through an ICL-enabled (MC-AE) IFL, packet loss would be seen and last for a while. PR1219569

- When MCNH (multicast composite nexthop) is used, packet loss might occur when multicast traffic enters the Packet Forwarding Engine and exits the Packet Forwarding Engine in a different FPC. PR1219962
- On MX Series routers with enhanced subscriber management, performing a config commit that changes any dynamic profile data after the system has booted might result in login and logout CPS(connections per second) 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 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 aggregated Ethernet child links was reloaded, traffic loss would be observed after the MPC boots up due to incorrect programming of aggregated Ethernet interface on its Packet Forwarding Engine. PR1220934
- When MX Series router has MACsec under security and the include-sci option is configured, although the interface where MACsec is configured receives traffic with IMIX packet sizes, framing errors might be reported in the interface statistics. PR1221099
- PPPoE/DHCP subscribers fail to bind due to ProcessPADIFailedUiflNotActive/SML_CLIENT_DELETE_SDB_ADD_FAILED errors after continuous login and logout, and subsequent login will fail. PR1221690
- Starting in Junos OS release15.1R1, the behavior of storage devices enumeration in kernel level has been changed. Device enumeration in legacy Junos OS releases (before 15.1R1) will show CF and Disk as ad0 and ad1 respectively Device enumeration in Junos OS Release 15.1R1 and later will show CF and Disk as ad1 and ad0 instead in the result of "show chassis hardware". This might be inconsistent for other result of output, such as "show system boot-messages" and "show log messages". PR1222330
- On setup with IRB configuration and non-enhanced-ip mode, when certain actions which result in the underlying aggregated Ethernet interface of IRB going down, the backup Routing Engine may experience 'panic' and hence reboot. The panic occurs because the backup Routing Engine cannot allocate the next-hop index requested by the master Routing Engine. Because the panic and reboot happen on backup Routing Engine, routing, forwarding, and any other functionality will not be affected. Some examples of triggers are continuous child link flaps of the aggregated Ethernet interface or back-to-back commits of different IRB configurations, and activating/deactivating the bridge family on an underlying interface. PR1222582
- In an enhanced subscriber management environment ("set system services subscriber-management enable") in which case the 'remove-when-no-subscribers' statement is configured in auto-configure stanza, when the last subscriber logs out (which triggers) dynamic VLAN IFL removal) and immediately then in close proximity a new subscriber logs in before the IFL is set to inactive, the dynamic profile deletion might be failed. Also, subsequent subscriber logins fail. This is a timing issue. PR1222829

- The "unnumbered-address" under dynamic profile is showing the wrong value. PR1222975
- The problem of tunnel stream getting misconfigured for LT interfaces is due to internal programming and the same has been corrected to evaluate multiple lt interfaces for FPC and PIC slot combination. PR1223087
- In MX Series Virtual Chassis with subscriber management environment, the bbe-smgd process may leak memory in the backup Routing Engine when running continuous subscriber login logout loop tests. It seems that memory utilization increases with each login/logout loop until it reaches 809 MB, and it does not increase beyond that. PR1223625
- In PPPoE subscriber scenario, after demux 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
- When you receive alignment errors on a 10 Gigabit Ethernet port, you may see a MAC control frames counter with a huge value. PR1224632
- SMID was coring when the query was thrown, because session database init was happened. A protection check introduced to check session database status.PR1225449
- The following error messages might be seen when there is traffic from subscribers with captive-portal-content-delivery service:cpcdd[29943]: %DAEMON-3: early: bad stored heap: heap-ptr=0x0 data-ptr=0x839f742 cpcdd[29943]: %DAEMON-3: opDel: bad stored heap: heap-ptr=0x1000000 data-ptr=0x839f0aa These error messages do not have any affect on functionality. PR1226782
- On MX Series routers, executing the command "show chassis ucode-rebalance" without a special FPC slot number, might cause chassisd to crash. PR1227445
- In a subscriber management environment, the 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 the time needed to complete execution of the routine exceeds. This message is harmless and can be ignored. However, sometimes time calculation yields incorrect results, and this issue has been corrected via this PR. PR1229967
- When adding or deleting a dynamic-tunnel destination network for IPv6 over IPv4 dynamic UDP tunnels, an rpd core file might be seen. PR1230152
- 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, the interface statistic remains unchanged after ISSU on MX Series Virtual Chassis platform. This issue in turn leads to the RADIUS volume accounting value remaining unchanged after ISSU. PR1230524

- The dynamic-profile service filter matches the traffic that is not defined in prefix-list applied to the filter. This causes the filter to not work not as expected or even match all the traffic. PR1230997
- 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) to appear as "*". PR1231868
- IPsec tunnels anchored on service-set are not cleared when ms interface inside IFL is disabled through CLI command. PR1232276
- Optional service session is terminating during session setup when optional service has configuration errors. PR1232287
- Some PFE statistics counters do not work in MPC7/8/9.

  1. Fabric input/output pps counters do not work in "show pfe statistics traffic"

  2. Output and fabric input/output counters do not work in "show pfe statistics traffic detail"PR1232540
- Packet Forwarding Engine statistics input packets pps counter may be inaccurate on MPC7E, MPC8E, and MPC9E. PR1232547
- Input framing errors increment on interfaces connected to MPC2E-NG with 4x10G MIC when interface is configured in "wan-phy" mode. PR1232618
- On XQ-based linecard, in a rare condition, when the FPC goes offline or online or when flapping occurs, some error messages might be seen. PR1232686
- Correct the value of module voltage, which was previously off 10 times, displayed in the interface diagnostics optics table for 2X100GE CFP2 OTN MIC. PR1233307
- High MPC5 CPU on a scaled setup with 64,000 - 128,000 subscribers due to XQ background service that collect internal statistics. PR1233452
- LSP-ping might fail and IP packets with options will not get mirrored in port-mirror environment. PR1234006
- For some SNMP traps the description does not match the event, for example: jnxTimingFaultLOESMCClear .1.3.6.1.4.1.2636.3.75.1.6 jnxTimingFaults 6 JUNIPER-TIMING-NOTFNS-MIB "A trap which signifies Loss of ESMC." PR1234083
- Due to a software bug, when an SFB goes offline/online, the HSLink crc error values are not cleared properly; this triggers an unexpected link error/ SFB check alarm for another SFB. PR1234224
- 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 an commit error. But there is no need to check when the Routing Engine is in backup state. PR1234453
- KRT queue is getting stuck happening because socket buffer is sending an incorrect value to the kernel and the kernel is returning error 'EINVAL -- Bad parameter in request'. PR1234579
- Phase jump is detected when using hybrid mode PTPoE with SyncE. PR1234685

- On MX Series routers with MPC7E/MPC8E/MPC9E, noise received on the console port might be interpreted as valid signals. This might cause login failure on the console port and login crash or even reloads. PR1234712
- When a session is started with a dynamic-profile service using the service volume, it is observed that volumes are checked every 10 minutes instead of every 5 minutes. PR1234887
- VLNS(VBNG) - Commit generated a "warning: requires 'l2tp-inline-lns' license" but a valid license is installed. PR1235697
- On MX Series routers, when per-packet load sharing is enabled under the aggregated Ethernet interface, egress traffic over the aggregated Ethernet interface might be dropped unexpectedly. PR1235866
- Junos Telemetry Interface authentication demon does not close the client connection properly keeping stale connections. Following command “show system connections | match JVISION_PORT” will show multiple stale connections. PR1235874
- The "show route forwarding-table all" command is needed for tlb (traffic load balancer) and srd (Service Redundancy Daemon) while these daemons are running. And these outputs are being collected from tlb script as well as srd script. The "show system commit" command is getting executed from default-junos-show script. When the CLI command is issued "request support information", "show route forwarding-table all" and "show system commit" are taken twice by RSI (Request Support Information). PR1236180
- On all platforms that support EVPN-VXLAN, the outer source MAC in the ARP reply packet header does not correspond to the inner virtual MAC if virtual MAC is configured. PR1236225
- When PIC-based MPLS J-Flow is configured and MPLS packets are being sampled at egress (to be sent to service pic), the sampled packets do not reach the service PIC, which results in no MPLS J-Flow flows getting created. PR1236892
- Due to a software bug, if there is an MPC6E slot#10 installed in an SFB2-based MX2020 router, and SFB#4~7 is offlined/online once, the next slot SFB will get 'SFB check alarm' unexpectedly. For instance, an SFB#4 offline/online triggers an SFB#5 check alarm. PR1237134
- In MX Series Virtual Chassis subscriber management environment, LI enabled DHCP subscribers may 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 if the DNS configuration is provided from the access-profile or pool. If such data is provided from RADIUS, the output is correct. PR1237525
- Due to lack of 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 4000 to 64,000. PR1237854

- When the interface configured under "router-advertisement" physically comes up for the first time, the rpd might repeatedly send the router-advertisement, which might result in as high as 100% Routing Engine CPU usage. PR1237894
- After the number of licenses for the scale-subscriber feature was exceeded, customer encountered endless logs on the backup Routing Engine every 10 seconds. PR1238615
- MPC9E may generate an FPC core file with Junos OS Release 16.1R2.11 when configured with "mixed-rate AE bundles" and "adaptive load balancing". The load-balancing techniques are orthogonal to each other. PR1238964
- MX Series router is sending accounting interim without the update-interval configuration statement. PR1239273
- In a BGP-PIC scenario, a change in the IGP topology (for example, a link failure in the IGP path) causes traffic outage for certain prefixes. This issue occurs because the unilist next hops for these prefixes are in a broken state. PR1239357
- Traceroute will not resolve VRF loopback address where SI and pseudointerface exist. PR1240221
- Subscriber Management: MIB ifJnxTable is not supported for subscriber interfaces. PR1240632
- Session database (SDB) 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
- During scaled subscriber setup, the lowest dynamic-profile CoS service rate might be applied to other sessions. PR1241201
- The PTP clock class changes are delayed. When PTP fails and the system goes into holdover, it will send clock class 6 for the next 10-15 minutes. When the system goes from holdover in state "locked". It will send clock class 248 for the next 10-15 minutes. PR1241211
- In some specific case, 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
- Auto route insertion (ARI) IPv6 routes installed for IPsec dynamic endpoints might disappear from the routing-table after performing a graceful Routing Engine switchover (GRES) with nonstop active routing (NSR) enabled. The issue is triggered for IPv6 ARI routes with masks of /98 or longer. PR1242503
- Currently MS-MIC supports a maximum of 2000000 routes scale. This includes all IPv4, IPv6, and MPLS routes in the system. When scale limit is exceeded, the FDB (forwarding database) memory will become exhausted and the MS-MIC will start to drop the routes and print logs. PR1243581
- On MX Series Virtual Chassis, some VBF flows are missing after FPC restart. PR1244832
- 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

- With gRPC subscription for telemetry data with 2 seconds frequency, the jsd process might crash. PR1247254
- When IGP/link flapping or running the **clear mpls lsp** command, because of the RSVP stale label entry, traffic for BGP prefixes that are pointing to LSP in inet.0/inet6.0 might get silently dropped or discarded. PR1247900
- SPMB reboot causes a fabric black hole that lasts for more than 1 minute in TXP-3D. PR1248063
- PADI dropped due to duplicate client. PR1248282
- The bbe-smgd process might crash in case of duplicate UID variable names. For example, all CoS configuration elements should be converted implicitly to internal variables so they can be automatically used for different purposes in the dynamic-profile configuration. The bbe-smgd process crash cannot impact the traffic flows for existing subscribers, but does impact the creation of new subscribers. PR1248725
- Only one IA-NA dhcpv6 (without 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 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 and this will allow only one session to be established even though CPE's did not request PD. PR1249837
- "JAM:PL: Registered attributes for %x \n" will be logged as INFO level. PR1250091
- MPC5E/MPC2E-NG/MPC3E-NG/MPC7/MPC8/MPC9 might crash in some cases due to a software defect. If queues associated with the L4 node get freed but the L4 node is not freed at that time, later when trying to free the L4 node, because the queues have already been deleted, then a NULL queue node will be received and the MPC crashes. PR1250335
- FPC ukern process might crash on Linux-based linecards (for example, MPC7/8/9 on MX Series) due to a bug related to ukern scheduler. PR1250691
- The 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 a denial-of-service attack. PR1250832
- Accounting statistics are not correctly preserved across ISSU upgrades. PR1250919
- On Junos OS platforms with rpd (routing protocol process), if some interfaces go down, which results in some peers going down or BGP-RR (route-reflector) re-advertising routes, rpd (routing-protocol daemon) process might crash. PR1250978
- During Routing Engine switchover or request MPC/MS-MIC online requests, character corruption is observed in the log. PR1251400
- There is an rpd problem sending route update messages to the kernel. The KRT queue used to send the messages can get into a state where no more messages can be sent to the kernel. This causes the RIB and FIB to get out of sync. This is a timing problem between multiple rpd threads. It infrequently occurs at very large scale. PR1251556

- When a MIC is re-inserted into the same slot, it is possible that the software may fail to read the software identification of the MIC. This results are misidentification of the MIC and not being able to initialize properly, and MIC0 info might disappear. It has no traffic impact. PR1252998
- If "indirect-next-hop-change-acknowledgements" is enabled, the rpd will request an acknowledgment from the kernel when creating the new forwarding next hop for the indirect next hop. In a rare scenario with multipath configured, the rpd might restart while waiting for an acknowledgment from the kernel and the deletion of the old forwarding next hop is queued. PR1254735
- On MX Series routers with MPC2E-NG/MPC3E-NG, the interfaces of these line cards might not come up when connecting to third-party transport switch. PR1254795
- In the output of 'show subscribers extensive' the first IP address from the Framed Prefix (returned in Framed-IPv6-Prefix) looks to be assigned to the subscriber interface although it is not. The fix removes incorrect data. PR1255029
- 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 due to an out-of-memory condition if there are continuous changes to multicast groups with IRBs over a long period of time. PR1255290
- In VMX platform, if a lot of subscribers login/logout activity occurs when there are a large number of flows (500,000), multiple riot core files might be seen. PR1255866
- Syslog messages may report "krt_decode_comp read a non specific nh from kernel nhid" This is a harmless debug message. PR1256197
- Adding an application set with inactive applications that are not defined under the [applications] hierarchy will lead to constant core files each time the service PIC boots back up. 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** after **commit full**. PR1258238
- In a subscriber service environment, the device control process (DCD) might restart unexpectedly during commit process after changes to ATM interface configuration is applied. PR1258744
- PPPoE subscribers are not coming up while verifying that IPCP renegotiation is happening properly for terminated PPPoE subscribers. PR1260836
- When using an AMS interface and running the show interfaces extensive command, the logical interfaces will show only 0 for the packet counters. PR1258946
- When TRI-RATE SFP-T is installed on MIC-3D-20GE-SFP-E, FPC will generate **HEAP: Free at interrupt level /Free interrupt violation!** syslog message when the interface is going down. PR1259757

- Due to a software bug, the QSFPP-40GBASE-LR4 (CLI name is QSFP+-40G-LR4) might remain down after fiber link flap. This issue is specific to this optics module. PR1259930
- Class of service (CoS) does not correctly classify egress L3 multicast traffic from an ingress VLAN bridge interface after a configuration change. PR1260413
- Only the first multicast IP packet was saved when waiting for a route to be resolved. This fix will save up to 20 additional IPv4 Multicast packets and send all saved packets after the route is resolved. PR1260729
- In MX Series BNG subscriber management environment, there could be a slight deviation in the service accounting statistics when the subscriber session terminated abruptly. PR1260898
- During multicast activation of dynamic subscribers via a service profile, the bbe-smgd daemon in the backup Routing Engine could sometimes crash. PR1261285
- In a subscriber management scenario, it is observed that an authenticated dynamic VLAN interface with an idle-timeout is removed if there are no subscribers on top and if "remove-when-no-subscribers" is configured at the auto-configure stanza. The dynamic VLAN interface should only be removed after its idle timeout expires if it stayed idle during this period. PR1262157
- There is a problem that MX Series routers use the wrong routing table to send out the ICMP network unreachable message back to the source; this might cause some problem on the end-user CPE. PR1263094
- Dynamic VLAN interface is logged out upon 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
- Currently when the CoS adjustment-control-profile (ACP) is configured with radius-coa using the adjust-less algorithm, cosd strictly follows the configured algorithm when (1) only service-profiles and/or CoA is used to apply rates to the subscriber flow and (2) no line rate adjustment protocols such as ANCP or protocol tags (for example, PPPoE-tags) are being used to apply updates. This results in undesirable complexity in applying service profiles in the order activated based on an ACP approach that is intended to control the comparison of a configured-rate and a line rate, where the former represents a policy and the latter the capabilities of the access loop. When only service profiles are in use, such that more than one service profile may be applied to the subscriber via RADIUS CoA and each service profile affects the shaping rate of the subscriber, the correct behavior is for CoS to ignore the algorithm when no line rate protocol is in use. Instead it should use a replacement semantic (logically the algorithm "adjust-always") to apply a service profile initiated via CoA in the order received. Thus a profile chain can be easily managed that includes the client profile and one or more service profiles, thereby allowing predictable and intuitive revert semantics during service-deactivation or re-activation scenarios. Once a line rate protocol such as ANCP is enabled and updates are received, only then should cosd follow the algorithm because it will then be performing comparisons with the configured rate and a line rate (where the intended goal is minimum (policy rate, line rate)). As a follow-on, the ACP

configuration syntax will be revisited because it is unnecessarily complex for the intended use case. PR1263337

- After router reboot or JSD process crash, sometimes the listening socket for JSD is not operational. PR1263748
- After running **show arp** with subscribers connected bbe-smgd can become unresponsive/slow to other CLI commands. PR1264038
- On MX Series routers with MPC7E/MPC8E/MPC9E installed, due to a race condition in reading optic state, after restarting MPC/MIC, extra link transitions might be seen during the period that the port is coming up. This is a timing issue and the affected port is random. The link might transform/flap multiple times before the link stabilizes. PR1264039
- On MX Series routers with MS-MPC, with the Ethernet frames with more than 2000 bytes of payload, the mspmand process that manages the multiservices PIC might crash. The traffic forwarding might be affected. PR1264712
- In some situations, MX Series LAC does not encapsulate packets received from CPE in l2tp tunnel if this subscriber has a static pp0 unit configured on the LAC side. This issue is causing a permanent traffic black hole for this subscriber and leads to PPP session flaps or in ability to establish a PPP session between CPE and LNS in case of using lcp re-negociation on the LNS side. PR1265414
- If the dynamic VLAN profile does not have IFF configuration (for example, family PPPoE or family inet), but has firewall filter configuration, firewall filter indixes will not be released after the dynamic VLAN is removed. This eventually leads to depletion of available firewall filter indixes. PR1265973
- Per IETF RFCs, IGMPv3 & MLDv2 reports not sent to IANA reserved multicast addresses 224.0.0.22(IGMP V3 ROUTERS) and ff02::16(MLD V2 ROUTERS) should be discarded. But BNG processes these reports. With this fix, the reports will be discarded and Rx error counter updated. PR1266309
- When VSTP is enabled on a double-tagged aggregated Ethernet logical interface and there is another single-tagged aggregated Ethernet logical interface configured with the same outer VLAN tag, then the incoming traffic on that VLAN is incorrectly hitting the AE_RESERVED_IFL_UNIT (AEx.32767) and the traffic is getting dropped. PR1267238
- It is possible to see a bbe-smgd core under certain boundary conditions on the standby Routing Engine with certain specific configurations. Because the core is on the standby no disruption in service is expected and the system recovers from this condition. PR1267646
- The CLI configuration command **set chassis effective-shaping-rate** is enabled for the MX104. PR1267829
- Command **show arp interface xe-x/x/x no-resolve | display xml** returns XNM errors in the ouput. PR1269170
- On MIC-3D-20GE-SFP-E and MIC-3D-20GE-SFP-EH, an interrupt threshold was introduced. If MIC error interrupts are more than the threshold (> 2500 per 5 min), the MIC will be restarted. Due to that change, MIC error interrupts will hog the CPU when restart is initiated. PR1270420

- In MX Series routers equipped with a next generation Routing Engine (RE-S-X6-64G and REMX2K-X8-64G), the following log messages might be displayed as error messages after a commit command is executed: **sdk-vmmd: %USER-3: is_platform_rainier:** Platform found as rainier. PR1271134
- The Routing Engine might stop all services after GRES or ISSU. This issue is caused by corrupted Berkeley DB file after GRES or ISSU.PR1271306
- Changing the mode of the interfaces causes the interface to go DOWN/UP. For the interface to be down, all the queues (in/out) associated need to emptied. Due to a certain condition, this is timing out, the queue is not getting emptied, and the interface pointer is not getting freed properly resulting in FPC crash. PR1273462
- On MX Series with MPC7E/MPC8E/MPC9E installed, if the ports on MPC that mix 10-Gigabit Ethernet (GE) and 40GE/100GE, after 40GE/100GE port is configured under an aggregated Ethernet bundle, some received packets might be incorrectly dropped. This is due to a misconfiguration on the Aggregated Ethernet MAC address under the Packet Forwarding Engine. This issue might happen after configuring 40GE/100GE as LAG member. PR1274073
- GRE MTU initialization: When GRE tunnels come up, the individual tunnel family MTU (which is V4/V6/L2 and so on) is updated based on its underlying interface family address MTU if the MTU is not configured exclusively under this GRE tunnel. However, Junos OS simply copies the MTU size, but it does not deduct the outer IP/GRE header length (20 + 4 bytes). The secondary issue is that while the underlying interface family address MTU size updates, the GRE tunnel MTU size will not be refreshed.

  PMTU discovery mechanism clarification:

  By default, the GRE tunnel source does not send any packets to discover PMTU. When traffic flows from the GRE tunnel source to the destination (or traffic entering GRE tunnel from outside) and if any intermediate router has a lower MTU and DF bit is set in the packet outer IP header, then that router sends an ICMP error message with error code 4 (indicates "packet too big" and cannot fragment because of DF bit) back to the GRE source router. If this ICMP message successfully reaches the source router, then the GRE interface MTU is updated with the MTU value suggested in this ICMP packet. After that, a timer is started in the GRE source router to keep this MTU value for this GRE tunnel within 5 minutes. After 5 minutes, the GRE MTU gets back its previous value, which is based on the underlying interface family address MTU or the configured MTU. However, during this 5-minute timeout, if another ICMP message is received with a lower MTU than the previously updated MTU (from 1st ICMP error packet), then GRE MTU is updated to reflect this new number and the timer is restarted. PR1274203
- Previous default behavior: when the bfd-admin-down under "routing-options static" stanza is not 'not'-configured, it was passive; that is, the static routes would not be deleted on bfd-admin-down. Now the default behavior is active, that is. static routes will be deleted on bfd-admin-down. PR1275973

High Availability (HA) and Resiliency

- On all platforms, if running ISSU, connection might be broken between the master Routing Engine and the backup Routing Engine. PR1234196

- With the local pp0 interface configured for IPv6 and router advertisement, if the other side of the interface is not configured for IPv6, rpd high CPU utilization might be seen. PR1243338
- Vmcores were generated on both VCMm and VCBm at the same time. PR1274438

Infrastructure

- The GNU debugger, gdb, can be exploited in a way that may allow execution of arbitrary unsigned binary applications. PR968335
- In an RSVP scenario, provision RSVP LSP with ldp-tunneling enabled and the LSPs configured with link protection, continuous kernel logs and LDP statistics timeout errors might be seen when executing **show ldp traffic-statistics**. PR1215452
- During the upgrade harmless "invalid SMART checksum logs" might be seen. This PR will suppress unnecessary "invalid SMART checksum logs". PR1222105
- 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
- If SSD contains a valid permanent (non-resettable) offline-uncorrectable-sectors positive value, smartd logs on the nonzero value by default every 30 minutes, which is too frequent logging considering that there has not been a change in the value. PR1233992
- 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
- Legacy Junos Kernel might generate a core file on userland_sysctl / sysctl_root / sysctl_kern_proc_env / panic_on_watchdog_timeout. PR1254742
- On Junos OS devices with legacy Free BSD (Free BSD version 6.X) based on Junos OS, the devices might crash and reboot if there is a defect in the Junos SDK based multi-threaded application that has been used. PR1259616

Interfaces and Chassis

- In MX Series Virtual Chassis setup, CFM sessions on aggregated Ethernet interface are not distributed to FPC when member-1 chassis are chosen as primary. PR1198447
- The **show interfaces terse routing-instance all** command has the wrong display format when there are multiple addresses. PR1207272
- If the configuration can be scaled to have the inner list to have more than 4000 VLANS, the commit VLAN configuration operations might fail. PR1207939
- The dcd cannot start after router reboot because of a non-existing IFL referenced in 'demux-options underlying-interface'. PR1216811
- MPC might crash during ISSU from Junos OS Release 15.1R1 to a later release when QSFP/CXP/CFP2 optics are present. PR1216924
- Previously the same IP address could be configured on different IFLs from different IFDs, but only in the same routing instance. Only one IFL 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

- PPPoE tunneled subscriber (L2TP) might get stuck in terminating state if radius sendsÂ Framed-IP-Address and Framed-IP-Netmask via access-accept in LAC. PR1228802
- When using the Ethernet OAM Connectivity Fault Management feature, if the remote end deactivates the "protocols oam ethernet connectivity-fault-management maintenance-domain" configuration, the interface will go down as expected. However, once the remote end activates the configuration, the local interface stays down. (The defect is introduced in Junos OS Release 15.1F5 branch and occurs in 15.1F5-S3 or later.) PR1231315
- When OAM CFM (connectivity-fault-management) MEP is configured on the LSI or tunnel interface that is on DPC card, every time a DMM (two-way frame delay measurement) or 1DM (one-way frame delay measurement) packet is received, certain harmless error messages might be seen. This is due to software time stamping not being used. The fix addresses the time stamp and suppresses the logs as well. PR1232352
- The configuration change in which 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". For example: user@router# show | compare [edit interfaces demux0 unit 7000 demux-options] - underlying-interface xe-0/1/0; + underlying-interface ge-0/3/9; user@router# commit re0: error: Bandwidth on IFL demux0.7000 cannot be greater than that of its IFD error: DCD Configuration check FAILED. error: configuration check-out failed. PR1232598
- There is no SNMP trap for dot1agCfmMepHighestPrDefect with value 0 reported when the OAM CFM session recovers from any other failed state. PR1232947
- On MX series platform acting as broadband network gateway (BNG), in Point-to-Point Protocol (PPP) scenario, when using the Internet Protocol Control Protocol (IPCP) or Internet Protocol version 6 Control Protocol (IPv6CP) for negotiation, if the router receives Configure-Request packet from the client, MX Series BNG sends the Configure-Request packet, but does not send the Configure-Ack packet (in case it does not receive the Configure-Ack that responds to the Configure-Request packet it sent). The behavior does not follow RFC 1661, which demands that both actions Send-Configure-Request (that is, ConfReq from MX Series to client) and Send-Configure-Ack (i.e. ConfAck from MX to client) must be conducted on the router without any significant delay. PR1234004
- On MX Series routers acting as broadband network gateway (BNG), in Point-to-Point Protocol (PPP) scenario the router can send LCP Terminate-Ack packet after 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

- Under a particular condition in configuring interfaces which have vlan-id/vlan-tags configured, the commit operation might fail with an error message. PR1234050
- T3 interface configured with "compatibility-mode digital-link" may fail to come up due to incorrect subrate. PR1238395
- If the MTU on BNG and CPE sides has different values, in a rare situation the MX Series router might calculate the MTU value for the corresponding pp0 IFL incorrectly. PR1240257
- When static PPP over Ethernet (PPPoE) subscriber is trying to negotiate a PPP session exactly at the time when Graceful Routing Engine Switchover (GRES) happens, the negotiation might fail and the following logs can be observed in the output of **show log message** command. **Jan 12 10:17:24.360130 allocateSession: IFL not available: pp0.1 1600!=1600** PR1245465
- In scaled subscriber management login/logout tests, jpppd might crash if the shmlog entries using the command **clear shmlog entries logname all** are cleared. PR1245848
- In some rare situations Ethernet Connectivity Fault Management Daemon (cfmd) might crash when committing a configuration where CFM filter refers to a firewall policy. When hitting this issue, all CFM enabled interfaces are down. PR1246822
- If more than one IFL (logical interface) is configured under the same IFD (physical interface), and VRRP is configured on one IFL without VLAN and the lower unit number IFL has a VLAN configuration, then vrrpd incorrectly carries the VLAN information from the lower unit number IFL to this IFL's 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 synchronized between the kernel and subscriber management demon. When using router advertisement on a static VLAN demux interface and not in a IP dynamic profile, a router solicit from customer equipment might not be answered by the MX Series router. This is dependant on which address the CPE is using. In this PR the option to configure the MX Series router to use EUI-64 address for the demux VLAN, will ensure that the addresses are synchronized between the demons. PR1250313
- On Junos OS platforms, cfmd process runs by default. When bridge-domain is configured, if performing a commit to configuration that related to physical interface/logical interface (IFD/IFL), cfmd memory leak might occur due to a software defect. As a result, the memory leak could cause cfmd crash. PR1255584
- The snmp-set command fails when the FPC/PIC/port has a value greater than 9 When the snmp-set command is issued, it encounters the following error due to incomplete port number in the command pushed. **Jan 18 10:49:53.626342 snmpd_process_nvset: talking to mgd (60001) Jan 18 10:49:53.626350 >>> xml to mgd >>> Jan 18 10:49:53.626418 RPC-REPLY ERROR: missing or invalid port number in 'et-10/0/' <<<<<<<<<< commit failed** PR1259155
- On MIC-3D-20GE-SFP-E or MIC-3D-20GE-SFP, when SFP diagnostic information is being read out periodically, due to misbehaving SFP or noise on the I2C BUS, SFP thread

might be hogging the CPU and a watchdog check will restart the MPC to recover. Enhancements will prevent the SFP thread hogging and MPC restart. PR1260517

- In a dual-stack PPPoE subscribers environment, when the PPP session has been in "OPEN" state, if the router receives a Conf-Request message from the client, it then sends a Term-Request message as a reply unexpectedly. PR1260829
- In a subscriber scenario, when traceoptions is enabled with flag GRES under PPPoE, if the subscriber username contains a format. (that is, the character "%") that cannot be successfully handled by the traceoption process, pppd might crash. PR1264000
- These types of messages might be observed with configuration changes in an MX Series Virtual Chassis environment: Mar 2 00:14:30 CHASSISD_IPC_WRITE_ERR_NULL_ARGS: FRU has no connection arguments fru_send_msg Global FPC 14 Mar 2 00:14:30 SCC fru_set_boolean: send: set_boolean_cmd Global FPC 14 setting hold-pic-online-for-fabric-ready on. These messages are benign. PR1264647
- In a PPPoE scenario, subscribers might get disconnected due to a keepalive failure when CPE is adding an additional data field in PPP Echo Request. PR1273083
- The message dot1agCfmMepHighestPrDefect might be reported in the SNMP trap with the value of -1 instead of 0 on recovery after RDI. PR1273278

Layer 2 Features

- When VPLS unicast traffic needs to be passed to a remote PE node via the LSI interface then go through the LAG interface to the L2TP network, packets could be dropped due to improper token handling. PR1240960
- In VPLS topologies the kernel may report the error "pointchange for TLV type 00000052 not supported on IFL <name> " in /var/log/messages where <name> is a VT or LSI interface used by VPLS. The trigger to cause the issue depends on timing and is most often seen with high VPLS pseudowrite scaling when multihoming is configured, but other triggers might apply as well. The problem might cause high RPD CPU utilization, which can slow routing convergence. PR1279192

Layer 2 Ethernet Services

- This issue occurs when running LACP between Juniper and Cisco devices with different timers (Juniper fast and Cisco slow) on both sides. On the Cisco side it take almost 90 sec to bring the interface down from the bundle. When one interface is removed from the LAG on the Juniper side, the lead on the Cisco side needs to time out to bring the interface down from the bundle. This results in unexpected outage behavior on the network. PR1169358
- If the DHCP relay in a forward-only routing instance receives an option-82 embedded DHCP discover, then session establishment might fail. This issue will happen only if forward-only is configured. PR1187766
- On MX Series routers, if chassis level configuration is used to offline the FPC after detecting major errors, the FPC will be offlined. But if the committing configuration is performed after offlining the FPC, the FPC will be brought online back again. PR1218304
- MX Series router is not including Delegated-IPv6-Prefix in accounting interim. PR1231665

- DHCPv6 renegotiation-lockout configuration command range has expanded to 4-600 seconds. This enables the customer to reduce the MX Series BNG wait time for responding to DHCPv6 solicit retransmissions messages according to their requirement. PR1234009
- 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 ISSU, which causes OSPF adjacency flapping. PR1240679
- In a large-scale unified ISSU testing, a MPC/FPC might go offline during the FRU upgrade phase of unified ISSU. PR1256940
- The duplicate-clients-in-subnet option82 feature has changed in the following way:

  When duplicate-clients-in-subnet option82 is configured, the client is identified using the circuit-id and/or remote-id of option82. Any other suboptions, for example, suboption 9 vendor specific, will not be used as a client identifier. Also, if duplicate-clients-in-subnet option82 is configured, existing clients will be identified using the circuit-id and/or remote-id of option 82 if available rather than being torn down. PR1257701
- During the DHCPv6 renegotiation lockout time, BNG does not accept any DHCP solicits with rapid commit options for further processing. This may slow down the subscriber initialization in relatively high packet drop access network segments. Fix for this PR eliminates the impact of DHCPv6 renegotiation lockout timer for DHCP solicits with rapid commit options PR1263156

MPLS

- When there are statically configured ingress and transit LSPs, due to a timing issue, there could be a scenario wherein the selfID used by the transit LSP might be allocated to the ingress LSP. Ingress static LSP does not reuse the same selfID during rpd restart, whereas the transit static LSP tries to reuse the same selfID. This leads to rpd crash due to the collision when the transit LSP tries to reuse the same selfID. PR1084736
- User is allowed to configure both "load-balance-label-capability" and "no-load-balance-label-capability" together. This is incorrect and confusing. PR1126439
- In some Inter-op scenario, sometime a new label is advertised without withdrawing the old label by peer. Under such scenario, Junos OS rejects 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. PR1168184
- If PCE-controlled LSP is enabled, when the command **no-install-to-address** is configured under PCE-controlled LSP, the command **no-install-to-address** might not be honoured due to a code issue. Routes might be installed for the destination of PCE-controlled LSP, which might not be desired when this issue happens. PR1169889
- When using RSVP-TE protocol to establish LSPs, make before break (MBB) might not be quit and will 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 MPLS OAM with mpls-tp-mode is enabled and the OAM failure-action is configured with make-before-break, the RSVP Explicit Route Objects (EROs) of new path might be removed after Make-Before-Break (MBB). The issue could be observed when BFD packets are dropped or the LSP path link goes down.PR1207039
- When dynamic-tunnel is configured but RSVP signaling is disabled, any configuration that affects dynamic-tunnels could cause the rpd process to crash. PR1213431
- Due to an imperfect fix for compatible issue between 64-bit RPD and 32-bit client applications (such as "mpls ping", "monitor label-switched-path", "monitor static-lsp", etc) on Junos OS Release 15.1F5-S3/15.1F6/14.2R7/15.1R4/16.1R1, the function of monitoring signaled or static LSP is broken on either 64-bit or 32-bit RPD. But the other 32-bit client applications (such as "mpls ping" etc) is not impacted. PR1213722
- In a scaled environment, when there are many unicast NHs related to the same transport LSP (for example, the same RSVP or LDP label), MPLS traffic statistics collection may 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 towards the backup Routing Engine and FPCs. PR1214961
- 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 and getting listed multiple times in the MPLS LSP. PR1222179
- Junos OS supports protocols mpls (MPLS) in the VRF routing-instance, but Junos OS does not support protocols connections (CCC) inside the VRF routing-instance. However, when ANY INTERFACE under protocols mpls (MPLS) inside VRF routing-instance is configured/added, then it affects protocols connections (CCC) inside Master/Main/Default Instance. For instances, if ANY CE FACING INTERFACE under protocols mpls (MPLS) in any VRF routing-instance is configured/added, it is deleting the data structure containing CCC information as Junos OS does not have CCC information inside the VRF routing-instance. PR1222570
- On MX Series routers with MPCs or MICs, if BGP-LU is configured with the entropy label. The entropy label value being generated might not provide a good load sharing result. PR1235258
- The rsvp-lsp-enh-lp-upstream-status is taking more time to synchronize on the backup Routing Engine on Egress side. PR1242324
- On MX Series routers, the LDP might fail to install LDP route in inet.3 table if IS-IS is configured with source-packet-routing and ldp-tunneling is enabled, which might cause the LDP to fail to install routes when IS-IS routes are present. PR1248336
- With nonstop active 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

- When multiple RSVP LSPs are in ECMP and configured with metric values, if one of the LSPs removed the metric, other LSPs in ECMP might not honor the configured metric. PR1261961
- During MBB (make-before-break), next-hop will change in Packet Forwarding Engine, RSVP route does not request a next-hop ACK before changing the route pointing to a new next-hop. When the scale is high, traffic loss can be seen for up to 1 second. PR1264089
- Label 0 is assigned as IPv6 explicit null label when "explicit-null" is configured for LDP. However, label 2 should be used instead of label 0. PR1264753
- With LDP session-protection configured, the LDP session for the remote LDP peer for rLFA (remote loop free alternate) might still remain up, even after rLFA is disabled or after the remote targeted LDP session is no longer needed by rLFA. PR1266802
- When a container LSP has >10 member LSPs, only the first 10 LSP will be shown in the **show mpls container-lsp name <lsp-name> statistics** output.PR1267774
- 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
- During LDP shutdown, route added and deleted by LDP in the inet.0 table may be in the process of being deleted but still in the inet.0 table. The **show route extensive** CLI command might cause RPD to crash when trying to display the task name for such LDP route. PR1272993

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. Kernel then creates a composite NH. In this case this appears to be P2MP MCNH, which gets created. When any member interface is not a Packet Forwarding Engine specific interface (e.g, Vt, LSI, IRB or any other pseudo interfaces), kernel throws this message indicating that FMBB cannot be supported. These messages are harmless and do not have any impact. PR1230465

Network Management and Monitoring

- MX Series BNG might send empty SNMPv3 responses for bulk-get requests to poll dot3adAggPortListPorts related OID's when using nondefault maxMsgSize settings. PR1207683
- In MX Series subscriber management environment, sometimes BNG responds to the SNMP get requests with "Error: status=5 / vb_index=0" for some of the interface related MIBs. PR1218206
- The statistics of OID ifOutError incorrectly includes ifOutDiscards, the buffer overruns are counted under ifOutErrors along with ifOutDiscards when SNMP Query is performed on ifOutErrors. PR1243071

- On all platforms, if changing the syslog configuration, the eventd process might stop sending syslog message to a configured syslog server. PR1246712
- SNMPv2 traps used to have the routing-instance information(context) in the community in the form context@community In SNMPv3, the same routing-instance information will be added to the contextName field of the SNMPv3 trap. For traps originating from a default routing instance, this field will be empty as it was earlier. PR1265288

Platform and Infrastructure

- NPC cored with reference to [0x41490f64 in trinity_policer_free (result_ptr=0x5d671f64, nh_ptr=0x5d671f78) at ../../../src/pfe/common/pfe-arch/trinity/applications/dfw/dfw_action.c:1049]. This type of NPC core can be observed with a dynamic configuration change to the policer. The processing time in attempting to update all associated policers was exceeded. PR1071040
- SNMP queries to retrieve jnxRpmResSumPercentLost will return the RPM/TWAMP probe loss percentage as an integer value, whereas the precise value (including decimal points) can be retrieved through the CLI by using the following commands: show services rpm probe-results and show services rpm twamp client probe-results. PR1104897
- In a CoS environment with shaping-rate configuration under interface, if flapping that CoS interface, the shaping-rate function does not take effect. PR1163147
- With the fix, XM-DDR3 boot diagnostics will return the test result of all XM-DDR3 components to the XM driver. If any XM-DDR3 component fails in the boot diagnostics test, the XM driver will abort the XM chip init process and report HW failure. The line card will not be brought up to online with any XM-DDR3 fail, causing a potential risk when sending corrupted packets to the remote Packet Forwarding Engines via the fabric streams. PR1166106
- When graceful Routing Engine switchover (GRES) is configured, the ksyncd crashes on the backup Routing Engine if a VPN static route has a network address as a next-hop. This occurs because the backup Routing Engine is not ready for a graceful switchover. PR1179192
- When multicast, vpls-flood or bridge-flood traffic, on an affected FPC type, with packet sizes ranging from 112 - 113 bytes or 108 - 109 bytes cross zone boundaries within the router (zones are defined below), traffic forwarding towards the fabric might stall. The following syslog entry will be reported "FO: Cell packing interface error". The MPC that reports this syslog error message needs to be restarted to recover from this condition. PR1180397
- IPv6 now defaults to a probe type of ICMP. Prior to this a probe type had to be explicitly specified. This change brings functional parity between IPv4 and IPv6 probe types with regard to a default probe. PR1183196
- Issue occurs if there is at least one python event-scripts configured with policy defended in configuration database. There are also some policies without the script action that hit the same warning. #commit full Jun 10 13:24:44 re0: [edit event-options] 'policy DOM-SIGNAL-CHECK' warning: Policy 'DOM-SIGNAL-CHECK' is defined in both Junos

OS configuration database and event script, ignoring the one defined in the event script. PR1190964

- In a very rare scenario, during a TAC accounting configuration change, the auditd daemon crashes due to a race condition between auditd and its sigalarm handler. PR1191527
- Insertion of an offlined MPC6E into the MX2000 chassis can cause the FPC temperature sensor to detect transient "WARM TEMP" condition, and the chassis FAN in the same zone goes to high speed. PR1193273
- Customer can now set the maximum datasize statement for JET scripts to up to 3 GB. PR1193948
- Interface link flaps could occur or MPC might generate a core file with any GRES on an MX Series Virtual Chassis. On an MX Series Virtual Chassis, MPC board selects a clock from the next reference after GRES,which is a line interface. If there is no signal on that line, then the clock is bad and link flaps could occur or the MPC might generate a core file. PR1194651
- On an MX Series router with an MQCHIP line card (MX Series routers with MPCs) with traffic-control-profile, if the overhead-accounting is configured with negative values, it might not work. The shape function will be affected. PR1195866
- junos:key attribute, which is emitted in the XML format of the configuration, will not be emitted in the JSON format of the configuration. PR1195928
- Blank firewall logs for IPv6 packets with next-header hop-by-hop is fixed. PR1201864
- On MX Series routers with MPC2 NG/MPC3 NG/MPC3/MPC4/MPC5/MPC6 installed, when configuring multiple lt interfaces with HQOS on a MPC, due to a software defect, when creating internal lt tunnel stream in Packet Forwarding Engine, the tunnel bandwidth will be overridden to max bandwidth(60G for MPC2 NG/3 NG, 100G for MPC/3/4/5/6). This causes all of the 256 internal FIFO resources to be allocated only two tunnels. The allocation for other tunnels fails due to lack of resources. As a result, only two lt interfaces can stay up, other lt interfaces will go down. PR1209065
- On MX2000, **show chassis hardware detail** might show MICs are installed even after MICs are removed. PR1216413
- MX Series routers with MPCs might crash after firewall filter configuration change is committed. PR1220185
- Routing protocol process (RPD) might restart unexpectedly if one of its TCP sockets is closed. PR1221183
- 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 process powers off the MPC via L2cs write the respective power registers, but in hardware it 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 secs. It occurs sometimes. PR1222071

- NTP peers failed to synchronize in symmetric active mode when there is significant downtime of one peer (for example, due to power maintenance, such as HW or SW upgrades). PR1222544
- IPv6 traffic learned on an L2/bridge/multilink interface and when it has been traversed through MPLS, core random packets might get classified incorrectly by the fabric, which leads to packet loss. PR1223566
- Interface firewall filters might get mixed up after Routing Engine mastership switchover with GRES disabled. PR1224995
- This is a race condition between database creation and database access. Rarely reproducible. There is no functional impact of the core. PR1225086
- Next hop used for Routing Engine generated TCP traffic might differ from the one used for Routing Engine-generated non-TCP traffic if the prefix not subjected to 'then load-balanced per-packet' action and is pointing to an indirect next-hop resolved via unilist next-hop (ECMP). Before the fix for PR1193697, this leads to non-TCP traffic generated from Routing Engine taking one unicast next-hop while TCP traffic generated from Routing Engine is load-balanced across different next-hops. After the fix for PR1193697 this behaviour might lead to non-TCP host outbound traffic taking one unicast next-hop, while TCP host outbound traffic takes another. PR1229409
- Firewall filter index mapping gets incorrect after Routing Engine switchover, due to the contents of "/var/etc/filters/filter-define.conf" getting wrongly changed after Routing Engine switchover. PR1230954
- The apply-path change bit does not seem to get applied when prefix-list is modified and the DFWD daemon, which waits for the policy-options, does not get notified and the apply-path function is broken. PR1232299
- In an AI-Scripts (Advanced Insight Scripts) environment, when there is 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
- Incoming interface index could not be used as a load balancing input factor under family multiservice if the traffic payload is a non-Ethernet frame. PR1232943
- FPC memory leak seen on T4000 FPC Type 5. PR1233003
- 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 a proper state transition on L2TP/LTS clients logout hence the license count was not getting updated. The fix will ensure the license count is updated on logout regardless of the daemon going through proper state transition or not. PR1233298
- NTP.org and FreeBSD have published security advisories for vulnerabilities resolved in ntpd (NTP daemon). Server-side vulnerabilities are only exploitable on systems where NTP server is enabled within the [edit system ntp] hierarchy level. A summary of the vulnerabilities that may impact Junos OS is in JSA10776. Refer to JSA10776 for more information. PR1234119
- Login for flow-tap DTCP-over-SSH service fails when SSH key-based authentication is configured for the flow-tap user. PR1234464

- In an MX2010/2020 environment with an MS-PIC with a J-Flow configuration, MX2010/2020 cannot sample multicast traffic when this multicast is copied to multiple interfaces. PR1237164
- FPC and Routing Engine might get stuck in high CPU when DDoS SCFD is turned on. PR1237486
- The auditd daemon is on LCC except SFC. So the auditd on LCC generates log message. PR1238002
- Due to a regression issue, the presence of errors or traps during ISSU might result in an LU/XL based FPC crash. PR1239304
- On MX Series routers with MPC5/MPC7/MPC8/MPC9, when a low value of temporal buffersize (for example,10k) is configured, the threshold in the drop rule in the Packet Forwarding Engine (PFE) differs from what is expected. PR1240756
- During an unified ISSU process, an MPC1E/2E/3E/4E or MPC-3D-16XGE-SFPP may restart unexpectedly. This issue shows up as an error in ppe_cfg_morph_ucode_instr( ) routine which can be seen in syslog messages. PR1241729
- For hardware platforms based on EA or XQ chips (such as MPC2E-3D-NG-Q), the minimum buffer value programmable in the Packet Forwarding Engine is modified from 4096 bytes to 1568 bytes. PR1246197
- An MPC/FPC may report LUCHIP EDMEM error during ISSU. This may cause inconsistency or incorrect forwarding information (FIB) inside the Packet Forwarding Engine. While the MPC is in the problem state, the Packet Forwarding Engine may experience packets lost. The issue should be self corrected after the ISSU process is complete and the Packet Forwarding Engine learns new FIB entries. However, if the problem persists, the MPC might need to be restarted. PR1249395
- The configuration database is locked when a user that was in "configure exclusive" is logged out unexpectedly. PR1250305
- When RADIUS accounting is configured, the Junos OS device will try for the maximum number of times when sending RADIUS accounting requests to a non-reachable RADIUS accounting server. When the last try is sending but the socket is closed due to the 'network is down' between Junos OS device and RADIUS accounting server, the auditd might crash. Auditd will get restarted automatically after it crashes. So accounting continues to work after auditd crashes. However, at the time of crash if there are some messages in the auditd queue that need to be sent out from Junos OS device to accounting server, those messages might get lost. After auditd gets restarted, the next event that has to be sent to RADIUS server, will be sent normally. PR1250525
- In a logical-systems environment, if there are some failures that cause Routing Engine switchover (not perform Routing Engine switchover manually), the Kernel routing table (KRT) queue might get stuck on the new master Routing Engine with the error "ENOENT -- Item not found". PR1254980
- On MX Series routers with MPC5E or MPC6E cards, 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. It might cause some unicast traffic over VPLS or BRIDGE to be dropped. PR1255073

- Packets are not encapsulated with GRE header after disable and reenable gr- interface and GRE tunnel traffic might get dropped. PR1255706
- During an unified ISSU, memory from the previous image related to hash tables is not properly recycled, which leads to blocks of physical memory being left unused. The crash is triggered by an attempt to create a memory pool using one of these blocks. PR1258795
- mgd might crash after executing the command **show ephemeral-configuration | display inheritance**. This option is unsupported. PR1258823
- If IX chipset-based mic(MIC-3D-20GE, for example) is used on an MPC that has two more mic slots, the **show pfe statistics traffic detail** command could display in/out pps statistics unexpectedly. PR1259427
- After an interface switch, when the MAC moves from one interface to another, the next hop is incorrectly following the MAC route, which has been corrected via code changes. PR1259551
- When a DHCP/BOOP reply packet is received from an unnumbered interface, the FUD process might fail. PR1260623
- After an ISSU upgrade, the WRED drop profile may not be programmed correctly, resulting in an incorrect WRED drop. PR1260951
- On an MQ chip-based MPC, some DDRIF checksum errors are observed, which might send traffic to a black hole. This PR also includes a chassis management alarm when there is a DDRIF checksum error on the MPC. PR1260983
- On an MX Series Virtual Chassis setup acting as an MVPN bud node and having a downstream local receiver and a PE node, traffic with few multicast groups are reported not being forwarded to the local receiver. PR1261172
- MX Series routers with FPCs might crash generating a core file when interface-specific firewall filters are configured with policers. PR1267908
- On all platforms, fast flapping of interfaces/fast changing of configurations might cause an RPD crash and BGP sessions will flap very quickly. PR1269116

Port Security

- The transmit delay interval is the maximum time the key server will wait before installing a new TX SAK (default value is 6 seconds). When MKA transmit interval is set to 6 seconds, during key roll over both transmit interval and delay interval timers expire at the same time and a new TX SAK gets installed on the key server before the RX SAK is installed on the peer node causing traffic drop. PR1257041

Routing Policy and Firewall Filters

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

Routing Protocols

- When a BGP speaker (router) has multiple peers configured in a BGP group, there is sometimes an inaccurate count of prefixes. This occurs when the BGP speaker receives a route from a peer and re-advertises the route to another peer within the same group. In such instances, the MIB object "jnxBgpM2PrefixOutPrefixes" for peers in the same group reports the total number of advertised prefixes in the group. MIB value "jnxBgpM2PrefixOutPrefixes" is defined as being used on a per-peer basis. However, it is instead being used to report prefixes on a per-group basis. To display an accurate number of advertised prefixes, use the **show bgp neighbor** command. PR1116382
- 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) may occur for every PIM leave. If the memory is exhausted, the rpd may 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
- In a BGP scenario with inet-mdt family configured under protocols BGP, route table <TABLE>.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 the master Routing Engine.PR1207988
- In large-scale BGP route environments with multipath configured, if BGP sessions go down simultaneously, the rpd might crash because it cannot finish multipath cleanup within a 10-minute limit. PR1209695
- If BGP and NSR are configured, then doing GRES might cause BGP to get stuck in NSR replication state. PR1210781
- 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

- BGP routes are rejected as cluster ID loop prevention check fails due to a misconfiguration. But when the misconfiguration is removed BGP routes are not refreshed. The fix for this issue sends a soft route refresh dynamically when a cluster ID is deleted. PR1211065
- When IS-IS is configured with overload timeout of 60 seconds and fragmented LSPs exist (for example, 25 IS-IS neighbors + 10K ipv4 routes + 1K ipv6 routes), if link flap/neighbor down/restart routing event is triggered, the IS-IS routes might miss in the routing table, which might cause some protocol sessions to go down and traffic loss. PR1213166
- When changing the RD for an existing VRF with established chatty MSDP sessions or deletion/deactivation of MSDP session in the configuration, the rpd process might crash, which leads to traffic disruption. PR1216078
- The routing protocol process (RPD) on a backup Routing Engine might restart unexpectedly in a large BGP NLRI environment. PR1220651
- In the rare scenario with a maximum number of routes in the BGP RIB_OUT table (for example: there are more than 700K BGP routes in route table), if flapping BGP protocol, it might cause the rpd process to crash. PR1222554
- According to the SR draft, the SR Capabilities sub-TLV must be propagated throughout the level and should not be advertised across level boundaries (the S bit in Router Capability Flag is set to 0). When IS-IS segment routing is configured, the S bit in Router Capability Flag is set to 1, which means the IS-IS Router CAPABILITY TLV must be flooded across the entire routing domain. Thus it leads to the IS-IS adjacency failure with other vendor devices. PR1223448
- When doing multiple back-to-back GRES switchovers the BGP peerings might drop after three or more switchovers. PR1224330
- On the Junos OS devices during graceful restart, the restarting node might send "End of RIB" maker too soon to its helper nodes, before the actual route updates are completed, causing traffic loss. PR1225868
- On all platforms, if MPLS goes down due to link flap or FPC reboot or restart, rpd core could be seen. PR1228388
- When first multicast packet gets fragments because of bigger in size, the receiver in the MVPN scenario does not receive all fragments. The fix for this PR will make sure to wait untill the last fragment of the PIM register packet is received at RP before processing the PIM resolve request. After last fragment of register packet is received, the PIM register state is created and the PIM resolve request is triggered to install a multicast route. So, all fragments of the register packet will get forwarded to the receiver. PR1229398
- Junos OS 15.1 and later releases might be impacted by the receipt of a crafted BGP UPDATE which can lead to an rpd (routing process daemon) 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 may not work for the OSPF route in case if - also a LFA protection is available - there is not ECMP to candidate PQ node - PQ node's router-id belongs to different area. PR1230322
- When a BGP peer goes down on the peer device, there might be a case of freeing the BGP session resources twice on the Junos OS devices and it can result in an rpd crash. This issue occurs when graceful restart is enabled on the peering device. PR1230556
- In a rare condition after a BGP session flaps, BGP updates might not be sent completely, resulting in BGP routes being shown in the advertising-protocol table on the local end but not shown in the receive-protocol table on 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 red zone check fails and at the same time attempting to free reference memory block. The failure is caused when the red zone check receives an address that is not the beginning of a memory block. PR1232742
- Juniper Networks implemented BGP4-MIB (including bgpPeerTable and bgpPeerState) per RFC 4273. When there is IPv6 BGP neighbor, Junos OS is unable to return a correct value for the BGP peer. This is caused because bgpPeerTable/bgpPeerEntry is indexed by bgpPeerRemoteAddr, which is syntaxed as IpAddress, a 32-bit integer. But the IPv6 address is 128 bits. This will cause Junos OS to return 0.0.0.0, which is considered an invalid peer. PR1233790
- With BGP ORR (optimal-route-reflection) configured, if IS-IS LSP has more than one fragment and the LSP is purged (for example, a topology change after a link flap), then an rpd crash might be seen. PR1235504
- When a rib-group is configured with a nonexistant routing-instance, after deleting rib-group and deactivating static flow route, a stale route might be present in inetflow.0 rib. It might affect traffic forwarding. PR1236636
- When there are different LSPs towards the same egress endpoint and they are up and advertised in IS-IS or ISIS TE shortcuts are configured, the active route is expected to use the LSPs as ECMP next hops in inet.0. If in addition, RSVP load-balance bandwidth is configured it would be expected that traffic is load balanced taking into consideration the LSP's bandwidth. The later was not happening and the traffic was load balanced equally across all ECMP LSPs, which should not have been the case. PR1237531
- A combination of next-hop-self, add-path, and per-prefix-label on a BGP-LU (label-unicast) RR can cause the wrong MPLS.0 routing/forwarding swap state to be installed. PR1238119
- When a Juniper Networks device is running protocol BGP, and policy configuration is modified, an assertion condition might be hit where the routing protocol process generates a core file. PR1239990
- When sham-link is configured, doing a series of configuration changes about sham-link might cause sham-link not to bring up. PR1240391
- In a PIM scenario with BSR configured, after deleting a static RP configuration from another router, then checking an RP table on a BSR router, there might be a stale

bootstrap RP entry (which is the static RP deleted from another router) in the RP table. PR1241835

- 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 into corresponding outputs for these protocols. PR1245105
- In BGP configuration, the static rt-constrain feature is configured but family route-target is not present on any BGP configuration, RPD might generate a core file. This is due to cleanup code attempting to free state that was not created since family route-target was not configured. PR1247625
- On all platforms, OSPF next hop might keep flapping between rLFA (remote LFA) and LFA when multi-area (PQ node sits in different area) rLFA along with policy is configured PR1248746
- Junos OS supports the mechanism to preserve BGP routing details for a longer period from a failed BGP peer than the duration for which such routing information is maintained using the BGP graceful restart functionality. But due to a software defect, the LLGR (Long-Lived Graceful Restart) feature not working between a Juniper Networks PE to other vendor's RR. PR1248823
- The configuration statement "learn-pim-router" is not working properly and as a result PIM hello packet will not be forwarded over pseudowire and multicast traffic will be dropped when the statement is configured under igmp-snooping protocol. PR1251439
- Routing protocol process (rpd) might restart unexpectedly with a reference to ioth_session_delete_internal ( ) routine. PR1261970
- On MX Series routers, if enabling IS-IS segment routing but certain interface is not enabled RSVP, then it might cause corrupted TLV 22 of IS-IS (the size of the value part of the TLV exceeds 255), and it might cause rpd to crash for parsing the LSP (labeled switchover path). PR1262612
- If vrf-table-label is configured in carrier of carriers VRF routing-instance and a direct interface route is advertised from the VRF towards a CE device as BGP-LU (BGP Labeled Unicast) route, the MPLS label entry for the direct route is permanently stuck in the kernel routing table (KRT) queue. PR1263291
- On MX Series router, when configuring import policy of IPv6 prefix with a IPv4 next hop for a BGP neighbor, the Rpd might crash continuously. The rpd crashing stops only after deletion of the policy. PR1265224
- After configuring "family inet unicast extended-nexthop", in the BGP open message sent to the peer, "Nexthop AFI=2" should be in the message instead of "Nexthop AFI=3". PR1272807

Services Applications

- When using NAT on the MX Series router, the FTP ALG fails to translate the PORT command when the FTP client uses Active Mode and requests AUTH(SSL-TLS) but the FTP server does not use AUTH. PR1194510
- Backup SDG reported memory-usage zone in RED, live PIC cores have been collected and PICs have been restarted. PR1202872
- IDP policy is trashing with the 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 PR1209351
- The kmd process might hog CPU when continuously 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
- Once you disable the stateful-high availability feature for an interface and then reenable it for the same interface and it comes up as backup, we might see some delay before it actually starts the session synchronizing. PR1214015
- L2TP subscribers on LNS might get stuck in Terminated state. PR1215941
- When BNG receives an ANCP Port Up message for tunneled subscriber and this message contains Actual Interleaving Delay Upstream and Maximum Interleaving Delay Downstream TLVs, then corresponding AVPs in the incoming-call request message will be corrupted. PR1234440
- On Layer 2 Tunneling Protocol (L2TP) access concentrator (LAC) router where Access Node Control Protocol (ANCP) is used for bandwidth adjustment, L2TP Connect Speed Update Notification (CSUN) message to L2TP network server (LNS) might be sent after a short delay after ANCP Port Up with updated access line parameters was received. This delay is caused by current interaction scheme between ANCP and the L2TP daemons and can last up to 5 seconds. In a production network scenario this delay shouldn't be visible as the L2TP daemon checks for state updates each time when there is an L2TP packet that has to be sent or received. PR1234674
- PPPoE - L2TP subscribers might get stuck in Terminating state in longevity login/logout test. PR1235996

- When the stateful firewall flows time out repeatedly, there can be performance degradation on the MS-DPC PIC. This will eventually lead to MS-DPC unable to scale to the peak flows that we allow. 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 clean up of all logged in L2TP subscribers and L2TP tunnels. If the **clear services l2tp tunnel** command is issued when the clean up has not been completed, it is possible that the tunnel will not be cleaned up properly and get stuck in "Terminating" state. PR1249768
- With MS-MIC/MS-MPC used for NAT service, when changing the source-address under a NAT rule term for a BASIC-NAT translation type, all future traffic hitting the NAT term will be dropped. PR1257801
- L2TP Congestion Window set to 128 instead of 1 when tunnel is created. PR1265001
- Apply-group configuration may cause KMD process crash during "commit check" process, which causes IPsec tunnel establishment failures. After this fix, apply-group can be used. PR1265404
- On MX Series routers, in rare cases, If chassis tunnel configuration and the flowtaplite configuration are changed in the same commit, kernel might crash. This is a timing issue and the probability of hitting this issue is low. If NSR/GRES is enabled in the device, the impact might be low that the kernel crashes. On the contrary, if NSR/GRES is not enabled in the device, traffic loss/routing protocol restart might be seen. PR1273357

Subscriber Access Management

- The auth request does not cause the router to send RADIUS REQUEST message, "Failed to queue the request, will be queued in authd internal queue" PR1178813
- In a subscriber management environment with two or more RADIUS servers connected to an MX Series router, syslog is not generated when the RADIUS server is marked dead.PR1207904
- If RADIUS returns Framed-route="0.0.0.0/0" to a subscriber terminated on a Junos OS platform, this subscriber cannot log in due to an authentication error. PR1208637
- A 3GPPP-SGSN-MCC-MNC svp with value "999999" will send in all CCR-GY requests.PR1233847
- On MX Series routers with subscriber management, the DHCPv6 solicit packets with IA_PD option from the subscriber are ignored if DHCPv6 server doesn't have prefix to allocate for this subscriber, which is incorrect behavior. According to the RFC standard, DHCPv6 server should reply to such packets using special Status Code: NoPrefixAvail (6), which should be included in Advertise/Reply in case if no delegated prefix is available. PR1234042
- On MX Series router with dual Routing Engines, after router the GRES, if user adds traceoptions filter during GRES not ready period, the authd process might crash. PR1234395

- Call rate performance may be impacted under heavy load if there are large numbers of small linked address pools due to a bug in the allocation traversal algorithm. PR1264052
- **show network-access aaa statistics radius detail** can display an incorrect number of messages to the RADIUS server in case configured RADIUS server's are continuously flapping. PR1267307
- In an MX Series BNG environment, it was noticed that the Show network-access requests pending count continues to increase even though there are no pending authentication requests. PR1267702
- During L2BSA subscriber stress test, some of subscribers may report invalid Event-Timestamp to RADIUS. PR1270162

User Interface and Configuration

- An rpd memory is increasing and cannot go back after an IS-IS interface flap. If this memory leak reaches a high level that impacts the route calculating, it might cause unexpected network issue. PR1243702
- Some configuration objects are not properly handled by "delta-export" (dexp). This leads to an omission of the section of the configuration. PR1245187

VPNs

- In MVPN SPT-only mode scenario, the first multicast packet is lost when the multicast source is directly connected to the PE. PR1204425
- In NG-MVPN scenario, when "forwarding-cache timeout never non-discard-entry-only" is configured for an MVPN instance, even though the cache lifetime is shown as forever in the output of CLI command **show multicast route instance X extensive**, the route disappears after 7-8 minutes. PR1212061
- On Junos OS platforms, only VPLS supports automatic-site-id. Configuring automatic-site-id under the L2VPN instance could cause an rpd core. The fix has now been provided to add a commit check to disallow configuring automatic-site-id under a L2VPN instance. With this fix, commit error will be thrown if the user tries 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. If the rpd memory is exhausted, recovery can be accomplished by restarting rpd. If nonstop routing (NSR) is configured, the master Routing Engine can be switched over to the standby Routing Engine, causing the master rpd to exit and restart and free the leaked memory. PR1220363
- In NGMVPN scenario with asm-override-ssm configuration statement for source specific multicast (SSM) group, if you issue the **clear pim join** command on the source PE, downstream interfaces get pruned causing the multicast flow to stop. If you issue **clear pim join** one more time then the issue is resolved. PR1232623

- 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. The issue does not exist if NSR is not enabled. PR1241801
- An rpd crash might be observed with a segmentation fault after applying an L2VPN configuration followed by the **ping mpls l2vpn** command. PR1272612

Resolved Issues: 16.2R1

Forwarding and Sampling

- Sampling Route-Record Daemon (SRRD) process 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 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 Junos OS Release 14.2R5.8 and Junos OS Release 14.2R6.5 severely reduce MX80 series available memory and therefore RIB/FIB scaling. PR1187721
- Starting with Junos OS Release 14.2R1, FPC offline could trigger Sampling Route Record (SRRD) daemon restart. PR1191010
- On MX Series 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 MX Series Virtual Chassis (MX-VC) environment, the private local next hops and routes pointing to private local next hops are sent to Packet Forwarding Engine from master Routing Engine and not sent to slave Routing Engine, then an 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 new master Routing Engine and sent to Packet Forwarding Engine, the Packet Forwarding Engine might crash due to a stale next hop. PR951420
- 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. PR 1042737 removes this restriction, and a ps interface may be anchored to an IT interface without explicit tunnel-bandwidth configured. PR1042737
- Wrong byte count was seen in the ipfix exported statistics packets for mpls flows . This issue is taken care now . PR1067084
- The configuration support for enabling ingress and egress layer2-overhead is available in dynamic-profile but the functionality is not supported in Junos OS Release 15.1R3 and Junos OS Release 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

- If any linecard crashes early during unified ISSU warmboot, the CLI might report unified 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
- During SIB yanking (pulling a SIB out without offline) on PTX platform with FPC3, it is possible that traffic may be dropped resulting in an overall reduction in traffic throughput. PR1162977
- On rare occasions the transport daemon may generate a core dump after a configuration change. PR1164377
- With 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
- Sampled continues logging events in traceoption file after traceoption for sampled deactivated. This can be hit if there is no configuration under 'forwarding-options sampling' but other configuration for sampled is present (for example, port-mirroring). PR1168666
- When MS-MPC is used, if any bridging domain related configuration exists (for example "family bridge", ""vlan-bridge"", ""family evpn", etc), in some cases, continuous MS-MPC crashes. Hence traffic loss may occur. PR1169508
- On MX Series with MS-MPC/MS-MIC, for some reason, out of order execution of instructions on MS-MPC/MS-MIC might happen and then causing the mspmand daemon (which controls the service pic and process the data) core and crash. PR1169946
- When a CFM down-mep is configured on a STP-blocked interface which is housed on a DPCE card, flooding of traffic in the local L2 broadcast network might happen, leading to side-effects such as flapping of OSPF sessions, BFD sessions, or similar. PR1174175
- 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 virtual tunnel hosted interface are on two different FPCs (Non-AE case). PR1174257
- MTU discovery may not be working due to lack of VRF info on egress card for BBE Subscriber traffic. PR1177381
- CGNAT-NAT64: Few port leak are observed for the EIM/EIF IPv4 traffic(2M sessions) from public side. PR1177679
- Changes are needed to support dedicated users for control and multicast traffic. This will avoid unicast traffic to be hashed to users doing ucode processing. On JUNOS OS

side, this PR introduces new CLI command **set chassis fpc X performance-mode num-of-ucode-workers Y**. PR1178811

- If "router-advertisement" protocol is configured in client ppp profile, unsolicited RA might be sent before the IPv6CP Configuration ACK is received. PR1179066
- A micro BFD session sourced from an interface's L3 address works even when the interface is not assigned the related UBFD address. PR1180109
- 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
- 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
- DA mac filter is missing on Child link of AE after FPC restart. PR1184310
- When IPv4 firewall filter have 2625/32 destination in prefix-list, filter attached to subscriber interface is found broken. PR1184543
- 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 ifd/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
- When ams-interface is configured in warm-standby mode without adding any members, configuration commit will lead to rdd core. PR1185702
- Next hop attribute in a framed route is not applicable anymore. Since subscriber IP address is used as the next hop in all cases, there is no need to have an additional attribute for next hop for framed routes. PR1186046
- Traffic destined to VRRP VIP address or transit traffic with destination mac as VRRP VMAC which has paylod beyond 166 bytes (excluding headers) are dropped as "my-mac check failed" on MPC7E/8E/9E. PR1186537
- After loading COS related configuration on MPC5E/MPC6E/MPC2E-NG/MPC3E-NG linecard, these error messages might be seen: "trinity_insert_ifl_channel:6449 ifl 495 chan_index 495 NOENT" "jnh_ifl_topo_handler_pfe(11591): ifl=495 err=1 updating channel table nexthop" PR1186645

- On MX Series 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
- On MX Series platform, while using routing-instance for EVPN, and traceoptions is configured under global "protocols evpn", configuration of "vtep-source-interface" under global "switch-options" would be rejected. PR1189235
- On MX240/MX480/MX960/MX2010/MX2020 platform, in rare cases, MPC4 line card might never come back online after rebooting the chassis by "request system reboot both-routing-engine" command. PR1190418
- 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 OS workaround does not work. PR1192914
- Configuring an RLT interface and rebooting the router shows the RLT interface down. The show l2circuit connection shows an mtu mismatch as the immediate cause. For example, the problem may be seen with the following configuration:

  show configuration interfaces rlt0 redundancy-group { member-interface lt-4/0/0; member-interface lt-4/2/0; } unit 0 { encapsulation vlan-ccc; vlan-id 600; peer-unit 1; family ccc; } unit 1 { encapsulation vlan; vlan-id 600; peer-unit 0; family inet { address 70.70.70.1/24; } } PR1192932
- With GRES (graceful-switchover) and nonstop-bridging configured in Juniper devices with dual Routing Engines, the backup Routing Engine might run into high CPU usage due to abnormally high CPU utilization by firewall daemon. The abnormally high CPU usage might impact the functions that backup Routing Engine works for. PR1193891
- 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
- 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

- 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 pfe 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
- 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
- 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

- On MX240/MX480/MX960 platform with RE-S-2000 Routing Engine, the Hard-Drive information on Routing Engine RE-S-2000 is missing in **show chassis hardware detail** output after upgrading to Junos OS Release 15.1 and later. This is just a display issue and this has no impact on any functionality. PR1205004
- J-UKERN.mpc0 core after filter configuration change on vMX. PR1205325
- This issue is identified as software defect and the fix is added in Junos Os Release 16.1R2 and above. PR1205914
- When PCEP is enabled and LSPs are undergoing changes, like make before break (MBB) for rerouting, the rpd has to send those updates to PCE. However, when the PCEP session to PCE goes down, these updates are cancelled, but the rpd fails to completely reclaim the memory allocated for these updates. This causes increasing 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
- Multicast traffic is incorrectly forwarded in the multicast vlan for a few seconds for multicast groups disallowed by Universal Call Admission Control policy PR1206598
- RLT interface configuration is not supported. PR1207982
- VC link "last flapped" timestamp is reset to "Never" on the new backup Routing Engine after MX VC global GRES switchover. PR1208294
- The cpcdd daemon might core and restart on the subscriber scenario with CPCD (captive-portal-content-delivery) service configured. PR1208577
- On MX Series platform 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, wrong destination MAC would be generated for the subscriber. PR1209430
- BGP PIC installs multiple MPLS LSP next hops as Active instead of Standby in Packet Forwarding Engine, this can cause a routing loop. PR1209907
- During GRES or unified ISSU, the BFD protocol state of a child ifd may not get replicated on the backup Routing Engine until bfd starts running on the new Active Routing Engine. PR1211015
- 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, it will be installed as an access-internal route instead of access route if it has /32 mask length. PR1211281
- Inline J-Flow - 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 therefor discard these VRRP packets. As a result, the VRRP status in both CE becomes Master/Master. PR1212796

- The MS-MPC/MS-MIC service cards might encounter a core when using certain ALGs or the EIM (Endpoint-independent mapping)/EIF (Endpoint independent filtering) feature due to a bad mapping in memory. PR1213161
- AE IFL targeted distribution feature now provides 4 level of prioritization. Please refer document attached in PR for more details. PR1214725
- Inline J-Flow service will not work after unified ISSU on MPC5E and above type line cards. PR1214842
- 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 that 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
- 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 Tomcat releases, parameterized family i-net 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 JUNOS OSnow supports extending the SSM groups defined in below CLI for dynamic subscribers using the BBE configuration: https://www.juniper.net/documentation/en_US/junos14.2/topics/reference/configuration-statement/ssm-groups-edit-routing-options.html PR1216515
- This issue happens only with RLT configuration and only on Junos OS Release 16.1 and beyond. PR1216991
- 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
- Continuous error messages are seen. PR1221340
- 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
- On rare occasions, offlining a MIC-3D-16CHE1-T1-CE MIC can cause a FPC core. This is very unlikely in general and chances of it happening are very low. There is no workaround for this except to upgrade to an image with this fix present. PR1223277
- On MX2020 router, when all the SFBs are yanked out, there is no available fabric in system, but FPCs remain online state. There is no problem in offlining these SFB/SFb2s. PR1227342

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

Infrastructure

- The issue is the gstatd process for 64 bit Junos image does not get to the correct path in the code and due to that gstatd process fails to start. PR1074084
- From Junos OS Release 15.1 and later, smartd 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

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
- The jpppd might crash with a core dump due to memory heap violation associated with processing MLPPP requests PR1187558
- 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. (for example, 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
- In very rare possibility, mpc can be crashed with coredump will be seen when cli command 'request chassis mic offline fpc-slot <fpc-slot> mic-slot <mic-slot>' is executed due to software bug that sfp diagnostics polling function tries to access already destroyed sfp data structure by mic-offline. With fix, software will check if sfp data is valid before tries. PR1204485

- If version-3 configuration statement is not configured, the command of "show vrrp detail|extensive|interface" display VRRP-Version as 2 for inet6 address family. The VRRP IPv6 never supported any VRRP version 2. It was always version 3. This issue is cosmetic but not actual impact on VRRP IPv6 functionality. The VRRP packets generated for i-net6 address family are of VRRP version 3. PR1206212
- 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
- If the configuration can be scaled to have inner list to have more than 4K vlans, the commit vlan configuration operations might fail. PR1207939
- 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 the jpppd process receives a reply message attribute from the radius or tacplus server with a character of %, it might cause the jpppd process to crash and cause the ppp user to be offline PR1216169
- On Junos OS Release 14.2 and later releases, if asymmetric-hold-time, delegate-processing and preempt hold-time is configured, when neighbor's interface comes up again, "asymmetric-hold-time" feature cannot be used as expected. PR1219757

Layer 2 Features

- A new static MAC is configured under AE interface, but the MAC of the LACP PDUs sent out is not changed. PR1204895
- In dhcp relay environment, when delay-authentication and proxy mode are configured at same time. Jdhcpd may core due to NULL session ID. PR1219958
- During 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 cause the second ISSU in backup Routing Engine to be aborted in daemon prepare phase. It will not proceed further. PR1225523

MPLS

- Multiple RLFA backup gateways (one using spring inner label and other using TLDP label) can get programmed if the given node is PQnode to another node in the network that does not use SPRING RLFA backup for its LDP route, resulting in ECMP among backup next hops. Semantically both gateways provide the same protection path and TLDP based gateway is coming in the way of checking sanity of SPRING backup path. PR1176489
- 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
- In L3vpn with chained-composite-next-hop scenario, when receiving a TTL expired packet, the device will transmit a ICMP error message in a MPLS header, but the route next-hop for this ICMP error packet is discard, so the one error message will be logged. PR1194446
- When ldp is deactivated, there may still be route entries left in the ldp shadow routing table. RPD will core due to stranded route entries in the ldp routing table. PR1196405
- If RSVP link-protection optimize-timer is enabled, rpd memory might leak in "TED cross-connect" when a bypass LSP is being optimized. PR1198775
- This behavior is 16.1 release specific. When an ingress side link failure and LSP uses bypass path, LSR(DUT) cannot send proper "RSVP RRO" even if egress side topology changes. Please refer the following example. --- example --- 1. This is initial state. LSP of RRO has Link A and B IP address. bypass bypass Link C Link D +--------------------+ +------------------+ | | | | [Ingress LER] [LSR] [Egress LER] | | | | +--------------------+ +------------------+ Link A Link B strict path strict path 2. Link A is down. LSP of RRO has Link B and C IP address because LSR sends out RSVP RESV including proper RRO to Ingress LER. bypass RSVP RESV bypass Link C <-----+ Link D +--------------------+ | +------------------+ | | | | | [Ingress LER] [LSR] [Egress LER] | | | | +--------- X --------+ +------------------+ Link A Link B strict path strict path 3. Link B is down. LSP of RRO has Link B and C IP address because LSR does not send out RSVP RESV including proper RRO to Ingress LER. (wrong) bypass RSVP RESV bypass Link C <-----+ Link D +--------------------+ | +------------------+ | | | | | [Ingress LER] [LSR] [Egress LER] | | | | +--------- X --------+ +-------- X -------+ Link A Link B strict path strict path PR1207862
- 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
- In VPLS environment, if delete the routing-instance, in 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

Network Management and Monitoring

- In some cases the output of a **show version detail** command may pause and take over one minute to finish. Note that trying to abort with control-c does not shorten the delay to regain the cli prompt. PR1196129
- 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 commands <>' stanza. The fix escapes any newline chars which mitigates the issue. PR1200820
- RLI-24802 introduced in 16.1R1 caused some issues with snmp get-bulk. These changes are reverted from 16.1R2 PR1209561
- The reason for this new PR (1227121) is because the fix for PR-1126532 was accidentally reverted while committing code under another PR-1209561. Hence, the external content for this PR is same as: https://gnats.juniper.net/web/default/1126432#external_tab PR1227121

Platform and Infrastructure

- **show interfaces mac-database mac-address** <mac-addr> <intf-name>does not display any mac-specific traffic statistics data on Stout Line cards and also VMX for mac-learning enabled interfces mapped to i-net family. PR1012046
- 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 Traps - 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
- Kernel might crash when deactivate or deleting a static route that is configured to point to an unnumbered interface-name as qualified-next-hop. PR1118681
- XPATH expressions evaluations for YANG keywords yang leaf-ref/must/when are disabled by default. It means, even though YANG configuration has leaf-ref/must/when expressions, these expressions will not get validated/evaluated. PR1119972
- 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
- The issue happens after GRES. If commit on the new master during the config sync from the old master, commit might fail. PR1179324
- 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

- A customer has reported that if you mistakenly configure a static flow route at the wrong hierarchy in the configuration of an MX80 or MX104 that a core dump occurs upon commit. This does not happen on other MX Series platforms. PR1187469
- When access accept response from radius server contains class attribute, .class file is created. Normally .class file gets deleted in success scenario after the user logs in and reads the attributes. However, in error scenarios where the login fails or login succeeds but fails to read the user attributes, .class file is not deleted. Due to this, .class files will remain in /tmp folder. As multiple .class files are stored in /tmp folder, /tmp folder is running out of inodes. PR1187477
- 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
- On Trio platform with network-services enhanced-ip mode, FPC CPU goes high for several minutes (30mins) when bulk (10K) mac/arp are learnt via lsi interfaces, which caused traffic interrupt. The issue can be seen with various triggers (e.g. mac flush, FPC reboot or link flap etc) . PR1192338
- Syslog storage in a file could abruptly stop due a race condition in handling log file rotation. The fix is available from Junos OS Release 16.1R2 and later.PR1195239
- When using delta-export , on commit full the configuration on backup Routing Engine will be corrupted. PR1199895
- 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 follows: 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
- From Junos OS Release 15.1F2/14.2R4, validating configuration fails if commit scripts are used during software upgrade. PR1204881
- If inline J-Flow is configured in scaled scenarios, inline J-Flow 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
- A fusion setup can experience a leak of NH memory when MAC moves result in updated next hops. You must restart the MPC to regain the memory. PR1208514

- Workaround : Deactivate and Activate Inline J-Flow sampling instance How to Avoid 1. Don't make any Inline J-Flow specific configuration changes when service is not in steady state 2. configuration changes should be done in two steps. a) First configure the J-Flow related configuration except the Flow Table size. b) Flow table size should be changed in a separate commit from the rest of the J-Flow configuration. PR1210899
- 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, etc). These files are copied even if there was no corresponding change in the configuration thus unnecessarily increasing commit time. PR1210986
- 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 it resolved. PR1212191
- On MX Series platforms installed 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
- On MX2000, 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 ABDxxx79 MPC6E 3D CPU REV 12 711-045719 ABDxxx35 RMPC PMB MIC 0 REV 14 750-049457 ABCxxx22 2X100GE CFP2 OTN >>>>>>>> No MIC inside MIC 1 REV 26 750-046532 ABCxxx53 24X10GE SFPP >>>>>>>>>>No MIC inside XLM 0 REV 13 711-046638 ABDxxx59 MPC6E XL XLM 1 REV 13 711-046638 ABDxxx87 MPC6E XL PR1216413
- This rmopd core was caused by the NULL pointer in SW function. PR1217140
- For Junos devices supporting FreeBSD10 and with Junos OS Release 16.1R2, 16.1x60-D30 or 16.1x60-D35, when ephemeral database is in use and "persist-groups-inheritance" configuration statement is configured, daemons (for example, bbe-smgd, l2ald, ccmd, dcd but not limited) might crash after deletion of configuration from either ephemeral database or normal static configuration database. PR1217362
- MX Series with MPC/MICs might crash after firewall configuration change is committed. PR1220185
- Under certain conditions sync-other-re editing configuration warning might be displayed after reboot: lab@mx> configure exclusive warning: uncommitted changes will be discarded on exit Entering configuration mode Users currently editing the configuration: sync-other-re (pid 9220) on since 2016-10-03 00:16:36 PDT, idle 2d 05:47 sync-other-re (pid 9282) on since 2016-10-03 00:16:40 PDT, idle 2d 05:47 sync-other-re (pid 9333) on since 2016-10-03 00:16:49 PDT, idle 2d 05:47 sync-other-re (pid 9383) on since 2016-10-03 00:16:59 PDT, idle 2d 05:46 sync-other-re (pid 9433) on since 2016-10-03 00:17:07 PDT, idle 2d 05:46 PR1221723
- Usage of malformed certificates (such as those missing newline characters) may result in rejection. The symptom would be messages such as: mgd: error: Unable to derive certificate from input. PR1223764

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 polices 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, for example, 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
- BGP Flowspec provides for a BGP Extended Community that served to redirect traffic to a Virtual Routing and Forwarding (VRF) instance that matched the flow specification's Network Layer Reachability Information (NLRI). After the fix of the PR, all Junos platforms can support the following Redirect Extended Communities: +--------+-------------------+-------------------------------------+ | type | extended community | encoding | +--------+-------------------+-------------------------------------+ | 0x8008 | redirect AS-2byte | 2-octet AS, 4-octet Value | | 0x8108 | redirect IPv4 | 4-octet IPv4 Address, 2-octet Value | | 0x8208 | redirect AS-4byte | 4-octet AS, 2-octet Value | +--------+-------------------+-------------------------------------+ Please refer to RFC7674 for more information. PR1219724

Routing Protocols

- 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 a bgp peer has a hold time of zero configured the peer will not reach establishment. PR1138690
- 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

- A route which has an IPv6 nexthop which is resolved recursively over other routes may fail to resolve successfully. This problem could happen because the route resolver may incorrectly use the IPv4 family resolution tree to resolve the nexthop rather than the correct IPv6 resolution tree. As a result, no route covering the IPv6 nexthop address can be located so the route with the IPv6 nexthop remains unresolved and unusable. PR1192591
- 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-MPVN 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
- On Juniper devices with BGP flowspec and Graceful-Restart for BGP configured, after the Routing Engine switchover, the firewall filter __flowspec_default_inet__ might be missed, causing BGP flowspec not working correctly. PR1213227
- When using 64-bit routing protocol process, if OSPF (either OSPFv2 or OSPFv3) is configured, the device may not handle the LS-Update correctly when receiving the max sequence number (0x7fffffff, which should not happen in normal course) and discarding it without acknowledging it as a newer copy in the database. The issue surfaced because a particular implementation was also setting the LSA-sequence number to max sequence number before flushing out the LSA which was not per RFC. PR1217373
- When a route in inet.3 has a conditional context associated with it (usually when conditional policy (policy with condition statement) applied on BGP), the rpd process might crash when IS-IS flooding LSP. PR1220533

Services Applications

- 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 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
- If radius Primary-WINS(Juniper-ERX-VSA) is set as 0.0.0.0, subscribers is rejected by Authd and does not 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

- If executing rpc get command without newline character at end of <rpc>, then it will cause script execution break for timeout of rpc-reply. PR1146379
- Configuration database is locked by "root" user when trying to commit vpls circuit configurations in "configure exclusive" mode. PR1208390
- If user enters configuration mode with **configure exclusive** command, after configuration is automatic rollback due to commit unconfirmed, 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 fails 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

- 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

See Also

- New and Changed Features on page 25
- Changes in Behavior and Syntax on page 43
- Known Behavior on page 52
- Known Issues on page 54
- Documentation Updates on page 115
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 124

Documentation Updates

This section lists the errata and changes in Junos OS Release 16.2R2 documentation for MX Series and T Series.

- Advanced Subscriber Management Provision Guide on page 115
- Subscriber Management Access Network Guide on page 115
- Subscriber Management Provisioning Guide on page 116

Advanced Subscriber Management Provision Guide

- The “Example: Configuring HTTP Redirect Services on the Routing Engine” topic shows an incorrectly formatted redirect URL, http://www.example.com*?=%dest-url%*. The correct format is http://www.example.com*/url=%dest-url%*.

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,” “Configuring Dynamic Access-Internal Routes for DHCP Subscriber Management,” “Configuring Dynamic Access-Internal Routes for PPP Subscriber Management,” “access (Dynamic Access Routes), “ and “access-internal (Dynamic Access-Internal Routes).”

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

- Support for the packet-triggered subscribers and policy control rule base (PTSP) feature was discontinued starting in Junos OS Release 13.1R1, but this was not reflected in the documentation. Text exclusive to PTSP has been removed from the *Broadband Subscriber Sessions Feature Guide*. This includes all CLI topics and the following chapters:
  - "Configuring the PTSP Feature to Support Dynamic Subscribers"
  - "Configuring the PTSP Partition to Connect to the External Policy Manager"
  - "Configuring PTSP Services and Rules"
  - "Monitoring and Managing Packet-Triggered Subscribers"

  Topics for other features that refer to PTSP are updated to report the end of support.
- The *Broadband Subscriber Sessions Feature Guide* did not report that you can suspend AAA accounting, establish a baseline of accounting statistics, and resume accounting. This feature was introduced in Junos OS Release 15.1R4.

  [See Suspending AAA Accounting and Baselining Accounting Statistics Overview.]
- Starting in Junos OS Release 15.1, the *Broadband Subscriber Sessions Feature Guide* and the CLI Explorer incorrectly included information about the **show extensible-subscriber-services accounting** command. This command is not present in the CLI. Instead, you can use accounting profiles to collect statistics from the Packet Forwarding Engine for Extensible Subscriber Services Manager (ESSM) subscribers. See Flat-File Accounting Overview for information about accounting for ESSM subscribers.

See Also

- New and Changed Features on page 25
- Changes in Behavior and Syntax on page 43
- Known Behavior on page 52
- Known Issues on page 54
- Resolved Issues on page 62
- Migration, Upgrade, and Downgrade Instructions on page 117
- Product Compatibility on page 124

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 MX 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.x. 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.x remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see Understanding Junos OS with Upgraded FreeBSD.

The following table shows detailed information about which Junos OS can be used on which products:

Platform	FreeBSD 6.x-based Junos OS	FreeBSD 10.x-based Junos OS
MX5, MX10, MX40, MX80, MX104, T4000, TX-MATRIX-PLUS	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 16.2

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 more information about the installation process, see Installation and Upgrade Guide and Upgrading Junos OS with Upgraded FreeBSD.

Procedure to Upgrade to FreeBSD 10.x based Junos OS

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

To download and install FreeBSD 10.x based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new **jinstall** package on the routing platform.

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

All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-16.2R2.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-16.2R2.9-signed.tgz
```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos package):

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-16.2R2.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-16.2R2.9-limited.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***

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.x) 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.x) 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 might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. For upgrading the host OS on these routers with VM Host support, use the junos-vmhost-install-x.tgz image and specify the name of the regular package in the request vmhost software add command. For more information, see the VM Host Installation topic in the Installation and Upgrade Guide.

NOTE: After you install a Junos OS Release 16.2 jinstall package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the request system software rollback command. Instead, you must issue the request system software add no-validate command and specify the jinstall package that corresponds to the previously installed software.

NOTE: Most of the existing request system commands are not supported on routers with the RE-MX-X6 and RE-MX-X8 Routing Engines. See the VM Host Software Administrative Commands in the Installation and Upgrade Guide.

Procedure to Upgrade to FreeBSD 6.x based Junos OS

Products impacted: MX5, MX10, MX40, MX80, MX104, T4000, TX-MATRIX-PLUS.

To download and install FreeBSD 6.x based Junos OS:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new **jinstall** package on the routing platform.

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

- All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

  ```
  user@host> request system software add validate reboot
  source/jinstall-ppc-16.2R2.9-signed.tgz
  ```

- Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (Limited encryption Junos OS package):

  ```
  user@host> request system software add validate reboot
  source/jinstall-ppc-16.2R2.9-limited-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***

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 might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 16.2 jinstall package, you cannot return to the previously installed software by issuing the request system software rollback command. 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 even from Junos OS Release 14.1 to Release 16.1. 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://www.juniper.net/support/eol/junos.html.

Upgrading a Router with Redundant Routing Engines

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

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

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

Upgrading Using Unified ISSU

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 16.2

To downgrade from Release 16.2 to another supported release, follow the procedure for upgrading, but replace the 16.2 package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the *Software Installation and Upgrade Guide*.

Changes Planned For Future Releases

- **Change in default behavior of traffic engineering shortcuts in labeled IS-IS segment routing**—In Junos OS Releases 15.1F6, 15.1F7, 16.2R1, and 17.1R1, traffic engineering shortcuts are enabled for labeled IS-IS segment routes, when you configure **shortcut** at the following hierarchy levels, so that both IS-IS and labeled IS-IS routes are populated in the routing table.
  - **[edit protocols is-is traffic-engineering family inet]** for IPv4 traffic.
  - **[edit protocols is-is traffic-engineering family inet6]** for IPv6 traffic.

  Starting with Junos OS Release 17.2R1 onwards, explicit configuration of traffic engineering shortcuts for labeled IS-IS segment routes is planned to be introduced by configuring **shortcuts** at the following hierarchy levels:
  - **[edit protocols is-is traffic-engineering family inet-mpls]** for IPv4 traffic.
  - **[edit protocols is-is traffic-engineering family inet6-mpls]** for IPv6 traffic.

See Also

- New and Changed Features on page 25
- Changes in Behavior and Syntax on page 43
- Known Behavior on page 52
- Known Issues on page 54
- Resolved Issues on page 62
- Documentation Updates on page 115
- Product Compatibility on page 124

Product Compatibility

- Hardware Compatibility on page 124

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 MX 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/.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the Hardware Compatibility tool.

See Also

- New and Changed Features on page 25
- Changes in Behavior and Syntax on page 43
- Known Behavior on page 52
- Known Issues on page 54
- Resolved Issues on page 62
- Documentation Updates on page 115
- Migration, Upgrade, and Downgrade Instructions on page 117

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 16.2R2 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/software/junos/.

NOTE: Unified ISSU is not supported in Junos OS Release 16.2R2.

New and Changed Features

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for PTX Series routers.

Release 16.2R2 New and Changed Features

There are no new features or enhancements to existing features for PTX Series in Junos OS Release 16.2R2.

Release 16.2R1 New and Changed Features

Hardware

- **The Interoperation for third-generation FPCs with first-generation and second-generation FPCs (PTX5000)**—Starting in Junos OS Release 16.2R1, when third-generation FPCs are installed on a chassis with first-generation and second-generations FPCs, the FPCs can interoperate with each other.

NOTE: For the third-generation FPCs to interoperate with the older generation FPCs, the **enhanced-mode** statement cannot be configured on the chassis.In such a scenario, the third-generation FPCs can provide the same functionality as the first-generation and second-generation FPCs. Any advanced features that third-generation FPCs might provide are disabled.

- **Support for P2-10G-40G-QSFPP and P2-100GE-OTN PICs on third-generation FPCs (PTX5000)**—Starting in Junos OS Release 16.2R1, the P2-10G-40G-QSFPP and P2-100GE-OTN PICs are supported on PTX Series routers that have third-generation FPCs installed.
- **New P3-15-U-QSFP28 PIC (PTX5000)**—Starting in Junos OS Release 16.2R1, the PIC P3-15-U-QSFP28 is supported on PTX5000 routers that have third-generation FPCs installed.

NOTE: To install the P3-15-U-QSFP28 PIC, you must have a third-generation FPC installed on your system.

Following is the available port configuration for each FPC:

- FPC3-PTX-U1-L and FPC3-PTX-U1-R—10 ports configurable as 10-Gigabit Ethernet ports (using a 4x breakout cable), 40-Gigabit Ethernet ports, or 100-Gigabit Ethernet ports.
- FPC3-PTX-U2-L and FPC3-PTX-U2-R—10 ports configurable as 10-Gigabit Ethernet ports (using a 4x breakout cable), 40-Gigabit Ethernet ports, or 100-Gigabit Ethernet ports.
- FPC3-PTX-U3-L and FPC3-PTX-U3-R—15 ports configurable as 10-Gigabit Ethernet ports (using a 4x breakout cable), 40-Gigabit Ethernet ports, or 100-Gigabit Ethernet ports.

High Availability and Resiliency

- **ISSU Feature Explorer**—The ISSU Feature Explorer is an interactive tool that you can use to verify your device's ISSU compatibility with different JUNOS releases.

  [See ISSU Feature Explorer.]

Interfaces and Chassis

- **Support for unicast RPF (PTX Series)**—Starting in Junos OS Release 16.2R1, you can configure unicast reverse path forwarding (RPF) to reduce the impact of denial-of-service (DoS) attacks on PTX Series routers that have third-generation FPCs installed.

NOTE: Unicast RPF is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Support for DCU and SCU accounting (PTX Series)**—Starting in Junos OS Release 16.2R1, destination class usage (DCU) and source class usage (SCU) accounting are supported on PTX Series routers that have third-generation FPCs installed.

  NOTE: DCU and SCU accounting are supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

MPLS

- **Support for IPv6 tunneling over an MPLS-based IPv4 network (PTX Series)**—Starting in Junos OS Release 16.2R1, IPv6 tunneling over an MPLS-based IPv4 network using IPv6 Provider Edge (6PE) is supported on PTX Series routers that have third-generation FPCs installed.

Network Management and Monitoring

- **Support for accounting profiles (PTX Series)**—Starting in Junos OS Release 16.2R1, you can configure accounting profiles to collect data on PTX Series routers that have third-generation FPCs installed.

  NOTE: Configuring accounting profiles is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

Routing Policy and Firewall Filters

- **Support for filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling (PTX Series with third-generation FPCs)**—Starting in Junos OS Release 16.2R1, filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling uses firewall filters to provide decapsulation of GRE traffic. The filter-based GRE decapsulation also supports routing-instance as an action.

  NOTE: Configuring filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

- **Support for the no-decrement-ttl tunneling attribute (PTX Series)**—Starting in Junos OS Release 16.2R1, you can configure the **no-decrement-ttl** tunneling attribute for filter-based generic routing encapsulation (GRE) for IPV4 and IPV6 tunneling.

  NOTE: The **no-decrement-ttl** tunneling attribute is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

VPNS

- **Support for Layer 3 VPN (PTX Series)**—Starting in Junos OS Release 16.2R1, Layer 3 VPN is supported on PTX Series routers that have third-generation FPCs installed.

NOTE: Layer 3 VPN is supported only when the **enhanced-mode** statement is configured at the **[edit chassis network-services]** hierarchy level.

See Also

- Changes in Behavior and Syntax on page 129
- Known Behavior on page 133
- Known Issues on page 134
- Resolved Issues on page 135
- Documentation Updates on page 142
- Migration, Upgrade, and Downgrade Instructions on page 142
- Product Compatibility on page 147

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 16.2R2 for the PTX Series.

General Routing

- **Support for deletion of static routes when the BFD session goes down (PTX Series)**—Starting with Junos OS 16.2R2, the default behavior of the static route at the **[edit routing-options static static-route bfd-admin-down]** hierarchy level is active. So, the static routes are deleted when the BFD receives a session down message.

Interfaces and Chassis

- **Message now displayed when SIB autohealing is complete (PTX3000 and PTX5000)**—Starting in Junos OS Release 16.2R2, the output of **show chassis fabric**

errors autoheal displays a message when SIB autohealing is complete, as shown in the following example:

```
user@device> show chassis fabric errors autoheal
Mar 30 01:43:00
Time                      Error log of first 100 errors
2016-03-29 23:46:23 PDT   Req: sib 0
2016-03-29 23:46:23 PDT   Action: SIB 0 (autohealing)
2016-03-29 23:54:52 PDT   Completed: SIB 0  (autoheal)
```

Management

- **Support for status deprecated statement in YANG modules (PTX Series)**—Starting with Junos OS Release 16.1R2, Juniper Networks YANG modules include the **status deprecated** statement to indicate configuration statements, commands, and options that are deprecated.

Network Management and Monitoring

- **Possible change to the object identifier (PTX Series)**—Starting in Junos OS Release 16.2R2, the many warnings that occurred previously during MIB loading in the SNMP client has been reduced. This update could change the resulting OID depending on the SNMP client-loading logic. For example, jnxProductACX1000 is defined under the two following nodes:

  - jnxProductACX1000 OBJECT IDENTIFIER ::= { jnxProductVariationACX 1 }
  - jnxProductACX1000 OBJECT IDENTIFIER ::= { jnxProductVariationACX1000 1 }

  Because the second definition is the duplicate, it is removed. Previously, the SNMP client referred to the second OID based on MIB loading logic, and then you would see a change in OID for the client.

- **Improved usage of wildcard in SNMP notify-filter OID (PTX Series)**—Starting in Junos OS Releases 13.3R10 and 16.2R2, the filter subtree using an asterisk (*) is correctly read as a wildcard and not as ASCII value of 42. The problem occurred in the following routers:

  - M Series running Junos OS Release 11.4R13.5, 13.3R7-S1
  - ACX2000 Series running Junos OS Release 12.3X54-D20.9
  - MX Series running Junos OS Release 14.1X50-D125

  A sample of the change appears in the output of the **show snmp v3** command:

  Old Output

```
Filter              Subtree            Filter       Storage      Status
name                                   type         type
nf1                 1.2.3.4.5          include      nonvolatile  active
  <<<<< Issue
nf1                 1.42.6             include      nonvolatile  active
```

```
 <<<< Issue
```

New Output

```
Filter              Subtree           Filter      Storage      Status
name                                  type        type
nf1                 1.*.*.4.5         include     nonvolatile  active
  <<< Fixed
nf1                 1.*.6             include     nonvolatile  active
   <<< Fixed
```

- **Update to SNMP support of apply-path statement (PTX Series)**—In Junos OS Release 16.2R2, SNMP implementation for the **apply-path** configuration statement supports only two lists:
  - **apply-path "policy-options prefix-list <list-name> <*>"**

    This configuration has been supported from day 1.
  - **apply-path "access radius-server <*>"**

    This configuration is supported as of this release.
- **Juniper MIBs Loading Errors Fixed (PTX Series)**—In junos OS Release 16.2R1, duplicated entries and errors while loading MIBs on ManageEngine MIB browser are fixed for the following MIB files:
  - jnx-gen-set.mib
  - jnx-ifotn.mib
  - jnx-optics.mib

  [See MIB Explorer.]
- **SNMP syslog messages changed (PTX Series)**—In Junos OS Release 16.2R2, two misleading SNMP syslog messages were rewritten to accurately describe the event:
  - OLD --**AgentX master agent failed to respond to ping. Attempting to re-register**
    NEW –- **AgentX master agent failed to respond to ping, triggering cleanup!**
  - OLD –- **NET-SNMP version %s AgentX subagent connected**
    NEW --- **NET-SNMP version %s AgentX subagent Open-Sent!**

  [See the MIB Explorer.]

Platform and Infrastructure

- **Improvements to MIB validation during Junos OS build (libjsnmp) (PTX Series)**—Many warnings can be ignored if produced while MIBs are compiling. The following warnings should be considered as errors because they can break the build:

  [0-9]:.*failed to locate—An OID failed to be located.

  [0-9]:.*redefinition of identifier—Redefinition of OIDs found in jnx-chas-defines.

  [0-9]:.*sequence-type-mismatch—Type mismatch found in sequence syntax of the table and actual OID type.

  [0-9]:.*cannot be imported from module—MIB failed to import because order is not defined properly.

Routing Protocols

- **Change in default behavior of router capability (MX Series)**—Starting in Junos OS Release 15.1F7, 16.1R4, 16.2R2, 16.1X65, and 17.1R1 and later, router capability TLV distribution flag (S-bit), which controls IS-IS advertisements, will be reset so that the segment routing capable sub-TLV is propagated throughout the IS-IS level and not advertised across IS-IS level boundaries.

System Logging

- **Changes in syslog MIB timeout messages (all routers and switches)**—Starting in Junos OS Release 16.2, timeout messages logged in **/var/log/messages** for requests sent from snmpd to mib2d have changed:.

  - The frequency of AgentX timeout logs has been reduced.

    When a request sent from snmpd to mib2d times out, a timeout message is put in the messages log. Previously, one timeout message was written in the messages log for each such request, leading to flooding the log file with repetitive messages. Starting in Junos OS Release 16.2, one syslog message is created for all requests that time out at the same instant.

  - The “clearing the current stats” part in the logs has been removed.

    The timeout message has been changed because the subagents stats on timeout are no longer cleared. For example, the following is an old log message:

    Jul 11 20:56:50 re1-bx04.cbf14 snmpd[2436]: %-LIBJSNMP_NS_LOG_WARNING: WARNING: AgentX session, /var/run/mib2d-11, noticed request timeout. Clearing the current stats. Request PDUs: 1482, Response PDUs: 471, Request variables: 37186, Response variables: 11866, Average response time: 1826.42, Maximum response time: 40162.90

    New messages read as follows:

    Jul 16 10:32:30 Gladiator snmpd[2241]: %-LIBJSNMP_NS_LOG_WARNING: WARNING: AgentX session, /var/run/mib2d-11, noticed 3 request timeout.

```
Jul 16 10:32:30 Gladiator snmpd[2241]: %-LIBJSNMP_NS_LOG_WARNING: WARNING:
AgentX session, /var/run/mib2d-11, noticed 1 request timeout.
Jul 16 10:32:30 Gladiator snmpd[2241]: %-LIBJSNMP_NS_LOG_WARNING: WARNING:
AgentX session, /var/run/mib2d-11, noticed 1 request timeout.
```

User Interfaces and Configuration

- **Integers in configuration data in JSON format are displayed without quotation marks (PTX Series)**—Starting in Junos OS Release 16.2R2, integers in Junos OS configuration data emitted in JavaScript Object Notation (JSON) format are not enclosed in quotation marks. Prior to Junos OS Release 16.2R2, integers in JSON configuration data were treated as strings and enclosed in quotation marks.
- **Changes to the XML and JSON output when displaying the differences between the candidate and active configurations (PTX Series)**—Starting in Junos OS Release 16.2R2, when you compare the candidate and active configurations and display the differences in XML or JSON format, for example by using the **show | compare | display (json | xml)** CLI command or the **<get-configuration compare="rollback" format="(json | xml)">** RPC, the device omits the **<configuration>** tag in the XML output and omits the **configuration** object in the JSON output if the comparison either returns no differences or if the comparison returns differences for only non-native configuration data, for example, configuration data associated with an OpenConfig data model.

See Also

- New and Changed Features on page 126
- Known Behavior on page 133
- Known Issues on page 134
- Resolved Issues on page 135
- Documentation Updates on page 142
- Migration, Upgrade, and Downgrade Instructions on page 142
- Product Compatibility on page 147

Known Behavior

There are no known limitations in Junos OS Release 16.2R2 for PTX 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 126
- Changes in Behavior and Syntax on page 129
- Known Issues on page 134
- Resolved Issues on page 135
- Documentation Updates on page 142

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 16.2R2 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

- PTX 100GbE-LR4 interfaces may 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 100GbE-LR4 interface uses the "line clock" sources and when it does not have any external clocks from BITS-a or BITS-b, offlining the PIC, which is recovering clock from line, brings the "line clock" down and the reference clock is switched from "line clock" to "holdover". This reference clock transition may cause a large clock phase-shift in the 100GbE-LR4 CFP modules, and this phase-shift may cause the output optical pulse waveform distortion on the 100GbE-LR4 interfaces. Hence, it results in interface flap. This issue cannot be fixed by software due to hardware limitation. PR1130403
- While upgrading from 15.1F based images to 16.x+ images or downgrading from 16.x+ images to 15.1F based images, if validate option is enabled, there may be a chassisd crash and upgrade/downgrade will fail. This issue should not be seen if both base and target images are from 15.1F train or 16.x+ train. PR1171652
- Major errors might be seen on MPC3/FPC3 with 1X100 and 5x100 DWDM MIC/PIC. user@router> show chassis alarms no-forwarding 1 alarms currently active Alarm time Class Description <timestamp> . Major FPC 3 Major Errors The following messages are seen in the logs: fpc3 Cmerror Op Sub Set: CORDOBA : CORDOBA(3/0) link 0 : DSP loss of lock fpc3 Cmerror Op Sub Set: CORDOBA : CORDOBA(3/0) link 0 : DFE tuning failed alarmd[16241]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 3 Major Errors craftd[15906]: Major alarm set, FPC 3 Major Errors. PR1212089

MPLS

- LDP to BGP stitching with eBGP indirect nexthop having implicit null label had never worked on PTX Series routers. It works only when BGP indirect nexthop has real label. Workaround (1) Ensure the peer advertises real label by adding another router between the egress and Ingress PE. (2) Use IBGP that gets resolved over LDP or RSVP-TE LSPs. This will ensure that the BGP indirect nexthop has real label. PR1254702

Platform and Infrastructure

- On PTX Series router, parity memory errors might happen in pre-classifier engines within a MPC. Packets will be silently discarded, as such errors are not reported and makes it harder to diagnose. After the change in this PR, CM-ERRORs, such as syslogs and alarms, will be raised when parity memory errors occur. PR1059137

See Also

- New and Changed Features on page 126
- Changes in Behavior and Syntax on page 129
- Known Behavior on page 133
- Resolved Issues on page 135
- Documentation Updates on page 142
- Migration, Upgrade, and Downgrade Instructions on page 142
- Product Compatibility on page 147

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.

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

- Resolved Issues: 16.2R2 on page 135
- Resolved Issues: 16.2R1 on page 139

Resolved Issues: 16.2R2

General Routing

- Under certain conditions in PTX Series platform with non NG-RE, the master Routing Engine may fail to relinquish the mastership upon Control Board (CB) internal switch failure. PR1132557
- On an FPC with Express chipset using multicast with an aggregate interface environment. When there are multiple network events, such as a multicast route churn or flap, especially when these changes are of multicast next-hop on an aggregate Ethernet (AE), the FPC may restart unexpectedly. PR1196644

- On PTX Series routers, when an ARP entry is learned through the aggregated Ethernet interface, and a route is pointing to that ARP nexthop, the ARP entry might not expire even though the ARP IP is no longer reachable. This issue is due to the route next hop on aggregated Ethernet interface getting stuck in unicast state even if the remote end is not reachable, and the RPD never gets to determine that ARP is invalid. The route next hop on the aggregated Ethernet interface should be shown in 'hold' state when the remote end is not reachable. PR1211757
- A higher number of QoS drops may be seen during large micro-bursts of traffic on the low priority queues on PTX Series FPC type 3 than FPC type 1 or 2. This is due to the delay band-width buffer being reduced to ~64ms on the PTX FPC type 3 versus ~200ms on PTX FPC type 1 or 2. PR1223440
- On PTX Series routers, if you repeatedly disable and enable a physical interface, memory leak of daemon l2ald (Layer 2 address learning daemon) might occur. PR1224271
- PTX Series FPC3 might receive noise on the FPC console port, and interpret it as valid signals. This might cause login fails on the console port, core files, or even reloads. PR1224820
- When a PTX Series router has inline J-Flow configured, and the interfaces where sampling is configured are receiving TCP traffic, "DMA - memory map failed" error messages might be reported by the FPC3-PTX-U2 card. PR1227687
- In a third generation PTX3000 system (PTX3K + FPC3 + SIB3), due to a timing issue, 10x100GE PIC may not come online upon a FPC restart and will remain offline with “FPC X PIC 0 Failure” message seen in alarms. PR1236048
- The "show route forwarding-table all" command is needed for tlb (Traffic load balancer) and srd (Service Redundancy Daemon) while these daemons are running. And these outputs are being collected from tlb script as well as srd script. The "show system commit" command is getting executed from default-junos-show script. When the CLI command "request support information" is issued, "show route forwarding-table all" and "show system commit" are taken twice by RSI (Request Support Information). PR1236180
- When the user configures a TPID value other than 0x8100 on a single tagged interface with the configuration command "vlan-tags outer <TPID>. <VLAN-ID>" , the TPID value 0x8100 will be used instead of the user-specified TPID value. PR1237687
- This issue occurs on PTX Series routers with a unilist application (such as ECMP) with member links with MPLS LSPs. If one member link has MPLS LSP protection configured and another does not, the router might send the wrong packets. PR1239634
- This is a day one behavior where there is traffic from multiple Ingress Packet Forwarding Engines to single Egress Packet Forwarding Engine, and the scheduler assign will have extreme rate distribution (very high to one queue compare to others), you see the high configured queue will not get all its shared of Tx-rate config. This is due to small size OQB on GS scheduler (an ASIC property).PR1241291
- Junos telemetry interface client might notice slower arrival of packets for the configured sensors. PR1243810

- On PTX Series routers, add 'set' parameter (optional) to CLI 'request system software add'. It provides a way to install multiple software packages and software add-on packages at the same time. PR1246675
- While processing lookup results, IRP block would raise an interrupt upon detecting an error condition. The interrupt would be active until the trapcode error is read. Under certain condition, software is not reading this trapcode error upon IRP interrupt. This issue causes the following syslog message to be generated:

  fpc5 INTR: throttle 60sec PECHIP[2]:pe.irp.intr.status:ap0_trap(0): (Count:1)

  fpc5 INTR: throttle 3630sec PECHIP[2]:pe.irp.intr.status:ap0_trap(0): (Count:3434)

  fpc5 INTR: throttle 3600sec PECHIP[2]:pe.irp.intr.status:ap0_trap(0): (Count:6841)PR1256736
- The following log message may be printed frequently during normal operation on the backup Routing Engine of a PTX5000 router when the Routing Engine type is RE-DUO-C2600:

  /kernel: mastership: sent other RE mastership loss signal PR1260884
- On PTX Series routers, if running interfaces on QSFP28 PIC in 10G mode, some of the interfaces on the QSFP28 PIC may not come up after a system reboot or PIC restart. PR1263413
- This only affects PTX5000 and PTX3000 platforms with third-generation FPCs. Software periodically monitors voltages on the FPCs to check if they are within the proper range. This change adjusts the expected values for voltages on certain power rails of the FPCs. In rare cases it is possible that a marginal FPC was operating inside the older limits but outside the new limits, in which case a new chassis alarm will be raised for that FPC. PR1263675
- If pio errors are encountered at a certain stage during asic initialization, some data structures are not correctly updated. Later an attempt to recover from the error automatically may detect data structure inconsistencies and abort, resulting in a ukern core file. PR1268253
- In PTX Series routers equipped with a next-generation Routing Engine (RE-S-X6-64G, REMX2K-X8-64G, RE-PTX-X8-64G/CB2-PTX), the following log messages might be displayed as error messages after a 'commit' command is executed: sdk-vmmd: %USER-3: is_platform_rainier: Platform found as rainier. PR1271134

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 an IFD and if guaranteed rate is not set for an IFL under this IFD, then the temporal buffer configured.

  The display of this error message is valid when the guaranteed rate is '0', but it is not valid when 'guaranteed rate' is disabled. PR1238719

Infrastructure

- In an RSVP scenario, provision RSVP LSP with ldp-tunneling enabled and the LSPs configured with link protection, continuous kernel logs and LDP statistics timeout errors might be seen when executing 'show ldp traffic-statistics.' PR1215452

Interfaces and Chassis

- Configuring ODU FRR under otn-options for 2x100G DWDM PIC is an unsupported command on the PTX Series router; incorrectly adding such a configuration could result in an FPC crash and restart. PR1038551
- When QSFP28-100GBASE-LR4/QSFP+-40G-LPBK PICs speed is configured at chassis hierarchy. DCD was not reading the speed specified in set chassis fpc <fpc> pic <pic> port <port> speed <speed> As a result, when IFDs created using this configuration are added in the aggregated Ethernet bundle along with IFD of any other kind of pics, DCD gave a commit error. DCD was able to read the speed for other IFDs in the aggregated Ethernet bundle and was not able to read the speed of IFDs on QSFP28 PIC. Therefore, the following speed mismatch commit error was seen: Interface ae0 with child links of mixed speed but link-speed mixed is not configured. PR1167780

MPLS

- When there are statically configured ingress and transit LSPs, due to timing issue, there could be a scenario wherein the selfID used by the transit LSP might be allocated to the ingress LSP. Ingress static LSP does not reuse the same selfID during rpd restart, whereas the transit static LSP tries to reuse the same selfID. This leads to rpd crash due to the collision when the transit LSP tries to reuse the same selfID. PR1084736
- After MBB, new LSP will not have "explicit route." PR1207039
- In an MPLS OAM environment, a rare timing condition can result in rpd crash when a memory clean task is delayed. PR1233042
- On PTX Series routers, the LDP may fail to install LDP route in inet.3 table if IS-IS is configured with source-packet-routing and ldp-tunneling is enabled, which might cause the LDP to fail to install routes when IS-IS routes are present. PR1248336
- When the rpd daemon is terminating, the process of signaling the deletion of all RSVP LSPs may take so long that a watchdog timer is firing, resulting in an rpd core. PR1257367

Platform and Infrastructure

- In a very rare scenario, during TAC accounting configuration change, the auditd process crashes due to a race condition between auditd and its sigalarm handler. PR1191527
- A race condition occurs between database creation and database access. This issue is reproducible, and there is no functional impact of the core. PR1225086
- The mgd might crash after executing the command "show ephemeral-configuration | display inheritance". This is an unsupported option. PR1258823

Routing Protocols

- 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 rare scenario with a maximum number of routes in the BGP RIB_OUT table (for example: if there are more than 700K BGP routes in route table), if flapping BGP protocol, it might cause the rpd process to crash. PR1222554
- On Junos OS devices during a graceful restart, the restarting node might send "End of RIB" maker too soon to its helper nodes, before the actual route updates are completed causing traffic loss. PR1225868
- On all platforms, if MPLS goes down due to link, FPC reboot, or restart, rpd core files could be seen. PR1228388

Resolved Issues: 16.2R1

Forwarding and Sampling

- Sampling Route-Record Daemon (SRRD) process 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 one certain family is not configured on all of the FPC clients (for example, 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

General Routing

- **OK** LED on the CB not lighting up in the following scenarios:
  - When you insert only one Routing Engine or CB on the router and power on . Only the **Master** LED lights up but not the **OK** LED
  - When you test the issue with either Routing Engine or CB on slot 0 or slot 1, leaving the other slot empty.
  - When both slots have the Routing Engine or CB inserted, the problem is not seen. Master and OK LED is lightup on the master CB, and the backup CB has the OK LED lit up.

Engineering is debugging this cosmetic issue. PR1115148

- The **clear services accounting flow** command should not be used in Junos OS Release 15.1F4 or Junos OS Release 15.1F5 on inline J-Flow on PTX5000 router for PTX Series. This command is specific to J-Flow and is not supported in these releases. PR1117181
- After booting up FPC3 on a PTX Series router, the internal link communication between some chips on the FPC might fail to establish and, as a result, the **Host Loopback Wedge** error message is displayed. PR1171101
- In a very rare cases, multiple Routing Engine switchovers may result in SNGPMB crash. PR1176094
- On a PTX Series router, with FPC3 card, after soft restart of SIBs (it could be GRES or performing "restart chassis-control immediately" if on the same Routing Engine), then offline/online of any SIB, traffic loss is observed. PR1177652
- For FPC3 on a PTX Series router, in rare scenarios, while restarting FPC, a PIC index mismatch issue might result in FPC crash if it is configured with inline-JFlow. PR1183215
- The FPC might generate a core file when issuing clear threads and show threads simultaneously. PR1184113
- SIB Link errors are seen during GRES, when mixed FPC types are present with EIP mode enabled. PR1192348
- On a PTX Series router, when Bits external clock is down/up, the incorrect SNMP trap, jnxFruRemoval(CB), is generated. jnxExtSrcLockAcquired should be 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
- On a PTX Series router with FPC3, if inline J-Flow is enabled with high scale of IPv4 and IPv6 routes and aggressive route flapping, a multiservice crash and FPC reboot might be triggered. PR1196793
- When inline sampling is configured in a PTX Series router with third-generation FPC, a debug messageis logged even though a debug command is not issued. PR1197695
- On a PTX Series router with FPC type 1 and FPC type 2, if there is a problem with the ASIC in the FPC, might cause FPC being disconnected from Routing Engine. PR1207153
- A vulnerability in IPv6 processing has been discovered that might 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. A malicious network-based packet flood, sourced from beyond the local broadcast domain, can cause the Routing Engine's 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 the legitimate ND packet times out. Refer to JSA10749 for more information. PR1207527

- On PTX Series routers, when an FPC type 1 or 2 is restarted, CoS profiles can be applied incorrectly to certain virtual output queues (VOQs). This can lead to RED drops on that VOQs for traffic those enters the router on the restarted FPC. PR1211509
- 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 of the**show chassis environment** command. However, the alarm for this condition is not raised under "show chassis alarms" if the alarm is to be raised during a system reboot. PR1216335
- Power budget values for PTX 5K chassis, FPC, and PICs have been revised. For routers operating on limited power, this can change the point where alarms for power-over-budget or insufficient power are raised or cleared. PR1216404
- The options accepted for "set chassis fpc <n> license-mode" configuration of PTX FPC3 are changed to "IR" and "R". PR1221096

Interfaces and Chassis

- If QSFP28-100GBASE-LR4/QSFP+-40G-LPBK PICs speed is configured at chassis hierarchy. DCD was not reading speed specified in (set chassis fpc <fpc> pic <pic> port <port> speed <speed>) and as a result, when IFDs created using this configuration are added in AE bundle along with IFD of any other kind of pics, DCD used to give commit error. DCD was able to read speed for other IFDs in AE bundle and was not able to read speed of IFDs on QSFP28 PIC and hence use to complain about speed mismatch Commit error: Interface ae0 with child links of mixed speed but link-speed mixed is not configured PR1167780

MPLS

- This behavior is Junos OS 16.1 release specific. When an ingress side link failure and LSP uses bypass path, LSR(DUT) cannot send proper "RSVP RRO" even if egress side topology changes. Please refer the following example. --- example --- 1. This is initial state. LSP of RRO has Link A and B IP address. bypass bypass Link C Link D +--------------------+ +------------------+ | | | [Ingress LER] [LSR] [Egress LER] | | | | +--------------------+ +------------------+ Link A Link B strict path strict path 2. Link A is down. LSP of RRO has Link B and C IP address because LSR sends out RSVP RESV including proper RRO to Ingress LER. bypass RSVP RESV bypass Link C <-----+ Link D +--------------------+ | +------------------+ | | | | [Ingress LER] [LSR] [Egress LER] | | | | +--------- X --------+ +------------------+ Link A Link B strict path strict path 3. Link B is down. LSP of RRO has Link B and C IP address because LSR does not send out RSVP RESV including proper RRO to Ingress LER. (wrong) bypass RSVP RESV bypass Link C <-----+ Link D +--------------------+ | +------------------+ | | | | [Ingress LER] [LSR] [Egress LER] | | | | +--------- X --------+ +-------- X -------+ Link A Link B strict path strict path PR1207862

Platform and Infrastructure

- 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

- On a PTX Series router with **chassis network-services enhanced-mode** configured, the default policy **junos-ptx-series-default** is not loaded correctly in case of some configuring operations, which can cause BGP routes not to be installed in the forwarding table as expected. To avoid this issue, reboot the router after any configuring operations on network-services. PR1204827
- MIB file was updated to use official names of released products only. No queryable objects were changed. PR1219906

User Interface and Configuration

- When **persist-groups-inheritance** is configured and you issue a rollback, you might see that the configuration is not propagated properly after a commit. PR1214743

See Also

- New and Changed Features on page 126
- Changes in Behavior and Syntax on page 129
- Known Behavior on page 133
- Known Issues on page 134
- Documentation Updates on page 142
- Migration, Upgrade, and Downgrade Instructions on page 142
- Product Compatibility on page 147

Documentation Updates

There are no errata or changes in Junos OS Release 16.2R2 documentation for PTX Series.

See Also

- New and Changed Features on page 126
- Changes in Behavior and Syntax on page 129
- Known Behavior on page 133
- Known Issues on page 134
- Resolved Issues on page 135
- Migration, Upgrade, and Downgrade Instructions on page 142
- Product Compatibility on page 147

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 143
- Upgrading a Router with Redundant Routing Engines on page 143

Upgrading Using Unified ISSU

CAUTION: **Unified ISSU is not supported in Junos OS Release 16.2R1.**

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 16.2

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.

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.

To download and install Junos OS Release 16.2R2:

1. Using a Web browser, navigate to the **All Junos Platforms** software download URL on the Juniper Networks webpage:

   https://www.juniper.net/support/downloads/

2. Select the name of the Junos OS platform for the software that you want to download.

3. Select the release number (the number of the software version that you want to download) from the **Release** drop-down list to the right of the Download Software page.

4. Select the **Software** tab.

5. In the **Install Package** section of the **Software** tab, select the software package for the release.

6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

9. Copy the software to the routing platform or to your internal software distribution site.

10. Install the new **jinstall** package on the router.

    

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

    All customers except the customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

    ```
    user@host> request system software add validate reboot
    source/junos-install-ptx-x86-64-16.2R2.9.tgz
    ```

Customers in the Eurasian Customs Union (currently composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package (limited encryption Junos OS package):

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-16.2R2.9-limited.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 might take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: You need to install the Junos OS software package and host software package on the routers with the RE-PTX-X8 Routing Engine. For upgrading the host OS on this router with VM Host support, use the junos-vmhost-install-x.tgz image and specify the name of the regular package in the **request vmhost software add** command. For more information, see the VM Host Installation topic in the Installation and Upgrade Guide.

NOTE: After you install a Junos OS Release 16.2 jinstall package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the jinstall package that corresponds to the previously installed software.

NOTE: Most of the existing **request system** commands are not supported on routers with RE-PTX-X8 Routing Engines. See the VM Host Software Administrative Commands in the Installation and Upgrade Guide.

Changes Planned For Future Releases

- **Change in default behavior of traffic engineering shortcuts in labeled IS-IS segment routing (MX and PTX Series)**—In Junos OS Releases 15.1F6, 15.1F7, 16.2R1, and 17.1R1, traffic engineering shortcuts are enabled for labeled IS-IS segment routes, when you configure **shortcuts** at the following hierarchy levels.
  - **[edit protocols is-is traffic-engineering family inet]** for IPv4 traffic.
  - **[edit protocols is-is traffic-engineering family inet6]** for IPv6 traffic.

  Starting with Junos OS Release 17.2R1 onwards, explicit configuration of traffic engineering shortcuts for labeled IS-IS segment routes is planned to be introduced by configuring **shortcuts** at the following hierarchy levels:
  - **[edit protocols is-is traffic-engineering family inet-mpls]** for IPv4 traffic.
  - **[edit protocols is-is traffic-engineering family inet6-mpls]** for IPv6 traffic.

See Also
- New and Changed Features on page 126
- Changes in Behavior and Syntax on page 129
- Known Behavior on page 133
- Known Issues on page 134
- Resolved Issues on page 135
- Documentation Updates on page 142
- Product Compatibility on page 147

Product Compatibility

- Hardware Compatibility on page 147

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/.

Hardware Compatibility Tool

For a hardware compatibility matrix for optical interfaces and transceivers supported across all platforms, see the Hardware Compatibility tool.

See Also
- New and Changed Features on page 126

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 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/.

Juniper Networks Content Explorer is a Web-based application that helps you explore Juniper Networks technical documentation by product, task, and software release, and download documentation in PDF format. Find Content Explorer at: https://www.juniper.net/documentation/content-applications/content-explorer/.

Documentation Feedback

We encourage you to provide feedback, comments, and suggestions so that we can improve the documentation. You can provide feedback by using either of the following methods:

- Online feedback rating system—On any page at the Juniper Networks Technical Documentation site at https://www.juniper.net/documentation/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at https://www.juniper.net/documentation/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/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://support.juniper.net/support/
- Search for known bugs: https://kb.juniper.net/
- 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://support.juniper.net/support/downloads/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://forums.juniper.net
- 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/.

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

18 July 2019—Revision 11, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

14 February 2019—Revision 10, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

7 February 2019—Revision 9, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

10 January 2019—Revision 8, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

7 June 2018—Revision 7, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

17 May 2018—Revision 6, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

12 April 2018—Revision 5, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

17 November 2017—Revision 4, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

27 July 2017—Revision 3, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

20 July 2017—Revision 2, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

13 July 2017—Revision 1, Junos OS Release 16.2R2– ACX Series, EX Series, MX Series, PTX Series, and T Series.

7 March 2017—Revision 6, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, andT Series.

12 January 2017—Revision 5, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, andT Series.

29 December 2016—Revision 4, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, andT Series.

13 December 2016—Revision 3, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, andT Series.

6 December 2016—Revision 2, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, andT Series.

29 November 2016—Revision 1, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, andT Series.