Release Notes: Junos® OS Release 16.2R1 for the ACX Series, EX Series, MX Series, PTX Series, T Series

14 February 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, T Series, and Junos Fusion.

These release notes accompany Junos OS Release 16.2R1 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 16.2R1 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

This section describes the features and enhancements in Junos OS Release 16.2R1 for ACX Series Universal Metro Routers.

Release 16.2R1 New and Changed Features

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

- **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.

- **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 {
  }
  }
  ```

- **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.

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.

See Also

- Changes in Default Behavior and Syntax on page 9
- Known Behavior on page 11
- Known Issues on page 11
- Resolved Issues on page 12
- Documentation Updates on page 12
- Migration, Upgrade, and Downgrade Instructions on page 13
- Product Compatibility on page 13

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.

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 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.

Network Management and Monitoring

- **Possible change is in the object identifier (ACX 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.

Platform and Infrastructure

- **Improvements to MIB validation during Junos build (libjsnmp) (ACX Series)**—There are many warnings which can be ignored if produced while MIBs are compiling. Following are some warnings that you need to consider as 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 in jnx-chas-defines.

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

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

See Also

- New and Changed Features on page 4
- Known Behavior on page 11
- Known Issues on page 11
- Resolved Issues on page 12
- Documentation Updates on page 12

Known Behavior

There are no known limitations in Junos OS Release 16.2R1 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 4
- Changes in Default Behavior and Syntax on page 9
- Known Issues on page 11
- Resolved Issues on page 12
- Documentation Updates on page 12
- Migration, Upgrade, and Downgrade Instructions on page 13
- Product Compatibility on page 13

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

- 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 4
- Changes in Default Behavior and Syntax on page 9
- Known Behavior on page 11
- Resolved Issues on page 12
- Documentation Updates on page 12
- Migration, Upgrade, and Downgrade Instructions on page 13
- Product Compatibility on page 13

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 4
- Changes in Default Behavior and Syntax on page 9
- Known Behavior on page 11
- Known Issues on page 11
- Documentation Updates on page 12
- Migration, Upgrade, and Downgrade Instructions on page 13
- Product Compatibility on page 13

Documentation Updates

See Also
- New and Changed Features on page 4
- Changes in Default Behavior and Syntax on page 9
- Known Behavior on page 11
- Known Issues on page 11
- Resolved Issues on page 12
- Migration, Upgrade, and Downgrade Instructions on page 13
- Product Compatibility on page 13

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 13

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 http://www.juniper.net/support/eol/junos.html.

See Also

- New and Changed Features on page 4
- Changes in Default Behavior and Syntax on page 9
- Known Behavior on page 11
- Known Issues on page 11
- Resolved Issues on page 12
- Documentation Updates on page 12
- Product Compatibility on page 13

Product Compatibility

- Hardware Compatibility on page 14

Hardware Compatibility

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

To determine the features supported on ACX Series routers in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at http://pathfinder.juniper.net/feature-explorer/.

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 4
- Changes in Default Behavior and Syntax on page 9
- Known Behavior on page 11
- Known Issues on page 11
- Resolved Issues on page 12
- Documentation Updates on page 12
- Migration, Upgrade, and Downgrade Instructions on page 13

Junos OS Release Notes for EX Series Switches

These release notes accompany Junos OS Release 16.2R1 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

- New and Changed Features on page 14
- Changes in Behavior and Syntax on page 15
- Known Behavior on page 15
- Known Issues on page 16
- Documentation Updates on page 17
- Migration, Upgrade, and Downgrade Instructions on page 17
- Product Compatibility on page 18

New and Changed Features

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

NOTE: The following EX Series switches are supported in Release 16.2R1: EX9200.

See Also
- Changes in Behavior and Syntax on page 15
- Known Behavior on page 15
- Known Issues on page 16
- Documentation Updates on page 17
- Migration, Upgrade, and Downgrade Instructions on page 17
- Product Compatibility on page 18

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 16.2R1 for the EX Series.

- Management on page 15

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.

See Also
- New and Changed Features on page 14
- Known Behavior on page 15
- Known Issues on page 16
- Documentation Updates on page 17
- Migration, Upgrade, and Downgrade Instructions on page 17
- Product Compatibility on page 18

Known Behavior

There are no known limitations for the EX Series switches in Junos OS Release 16.2R1.

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 14
- Changes in Behavior and Syntax on page 15
- Known Issues on page 16

- Documentation Updates on page 17
- Migration, Upgrade, and Downgrade Instructions on page 17
- Product Compatibility on page 18

Known Issues

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

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

- Authentication and Access Control on page 16
- Network Management on page 16
- Platform and Infrastructure on page 16

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

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

See Also

- New and Changed Features on page 14
- Changes in Behavior and Syntax on page 15
- Known Behavior on page 15
- Documentation Updates on page 17
- Migration, Upgrade, and Downgrade Instructions on page 17
- Product Compatibility on page 18

Documentation Updates

There are no errata or changes in Junos OS Release 16.2R1 for the EX Series switches documentation.

See Also

- New and Changed Features on page 14
- Changes in Behavior and Syntax on page 15
- Known Behavior on page 15
- Known Issues on page 16
- Migration, Upgrade, and Downgrade Instructions on page 17
- Product Compatibility on page 18

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 17

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

See Also

- New and Changed Features on page 14
- Changes in Behavior and Syntax on page 15

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 http://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

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

These release notes accompany Junos OS Release 16.2R1 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 http://www.juniper.net/techpubs/software/junos/.

CAUTION:

When In-Service Software Upgrade (ISSU) is performed on MX Series routers from Junos OS 14.2, 15.1, or 16.1 to Junos OS 16.2, the following MPCs and its MICs of both Ethernet and Non-Ethernet variations fail to come online after the ISSU process is complete:

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

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 16.2R1 for the MX Series and T Series.

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 < *op* >]

- **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.

- **IS-IS Purge Originator Identification TLV (MX Series)** Beginning with Junos OS Release 15.1 F4 and 16.2R1, Junos OS supports RFC 6232, *Purge Originator Identification TLV for IS-IS* , which defines a type, length, and value (TLV) for identifying the origin of a purge initiated by the IS-IS protocol. You can configure this feature to add this TLV to a purge along with the system ID of the Intermediate System (IS) that has initiated this purge. This makes it easier to locate the origin of the purge and its cause. A new show command **show isis purge log** is introduced to view the purge history and to identify the purge originator.

  [See IS-IS Purge Originator Identification Overview.]

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 35

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 16.2R1 for MX Series and T Series.

Management

- **Support for status deprecated statement in YANG modules (MX Series and T Series)**—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.
- **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.

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.

Operation, Administration, and Maintenance (OAM)

- **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)**—There are many warnings which can be ignored if produced while MIBs are compiling. Following are some warnings that you need to consider as 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 in jnx-chas-defines.

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

  [0-9]:.*cannot be imported from module—MIB failed to import due to order 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.]

Services Applications

- **Change in option name to configure inactive timeout for IKE ALG child sessions (MX Series)**—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**.

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.

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.

See Also

- New and Changed Features on page 19
- Known Behavior on page 38
- Known Issues on page 38
- Resolved Issues on page 42
- Documentation Updates on page 58
- Migration, Upgrade, and Downgrade Instructions on page 59
- Product Compatibility on page 68

Known Behavior

There are no known limitations in Junos OS Release 16.2R1 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.

See Also

- New and Changed Features on page 19
- Changes in Behavior and Syntax on page 35
- Known Issues on page 38
- Resolved Issues on page 42
- Documentation Updates on page 58
- Migration, Upgrade, and Downgrade Instructions on page 59
- Product Compatibility on page 68

Known Issues

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

- In certain scenarios, MX Series routers can periodically generate "libstats_del_stats_record: failed to delete the record from DB, db_ret: -30989, ret:-7" log messages. These messages are in fact harmless debug messages, but they were mistakenly categorized as error messages. In newer Junos OS Releases with the fix, these messages are properly categorized as debug messages. PR1186387

General Routing

- On MX5/10/40/80, the router may restart when the Routing Engine (RE) memory utilization is high, e.g. 95%. This restart can manifest in two ways:

  1. A kernel core is generated after restart;

  2. The watchdog restart is triggered and no kernel core is generated;

  In both cases the system will restart. PR1099998

- The jsscd might crash in static-subscribers scaling environment (for example, 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

- As described in RFC7130, when LACP is used and considers the member link to be ready to forward traffic, the member link MUST NOT be used by the load balancer until all the micro-BFD sessions of the particular member link are in Up state. PR1192161

- l2ald might thrash if targeted-broadcast is configured on EVPN irb. PR1206979

- Major errors might be seen on MPC3/FPC3 with 1X100 and 5x100 DWDM MIC/PIC.

  root@abc-r1-re0> show chassis alarms no-forwarding

  1 alarms currently active

Alarm time Class Description

2016-01-1 11:22:44 UTC 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

- There is no unified ISSU from Junos OS Release 15.1 and older releases to Junos OS Release 16.2R1. PR1222540

Interfaces and Chassis

- The command **show interfaces terse routing-instance all** has wrong display format when there are multiple addresses. PR1207272
- CFM sessions does not come up sometimes after chassis control restart is done. The WA is to deactivate the cfm session and activate it back. PR1233712

Layer 2 Features

- When "input-vlan-map" with "push" operation is enabled for dual-tagged interfaces in "enhanced-ip" mode, there is a probability that the broadcast, unknown unicast, and multicast (BUM) traffic may be blackholed on some of the child interfaces of the egress Aggregated Ethernet (AE) interfaces or on some of the equal-cost multi-path (ECMP) core-links. PR1078617

Layer 2 Ethernet Services

- If a client sends a DHCP request packet, and Option 55 includes PAD option (0), a DHCP ACK will not be sent back to the client. PR1201413

MPLS

- In scaled environment, when there are many Unicast NHs which are related to the same transport LSP (for example, 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

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

- In the dual Routing Engine' s scenario, the backup Routing Engine does not sync up the configuration change while deleting an inactivated interface from the master. So after the operation, the inactivated interface still exists on the backup Routing Engine. PR991081
- The mustd daemon might crash when large configurations are committed. PR1186326
- Multicast traffic might get dropped when the STP port role is changed. Work around is to toggle the IGMP Snooping membership. PR1193325
- After graceful switch-over on a certain scaled scenario, (for example, 4k IFLs on each PE port along with 1M+ BGP inet0 routes and 15k/10k ipv4/ipv6 IGP routes and GRES/NSR active) kernel may throw ?rnh lookup failed? error. In such case, ksyncd may also crash and may trigger a live kernel core on both Routing Engines. PR1217292

Routing Protocols

- If there are more than 500 AS numbers in the AS path the routing process could restart. PR461329

- In rare cases, rpd may write a core file with signature "rt_notbest_sanity: Path selection failure on ..." The core is 'soft', which means there should be no impact to traffic or routing protocols. PR946415
- 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

Services Applications

- When MS-PIC is running on T640/T1600/T4000, the number of maximum service sets is wrongly limited to 4000, instead of 12000. This might impact in scaled service (IPsec, IDS, NAT, Stateful firewall filter, etc) environment. PR1195088

Subscriber Access Management

- On MX Series platforms, when using RADIUS dynamic requests for subscriber access management, if the device detects that the CoA-Request it received is same with the one in processing progress, the router would send CoA-NAK packet back to the RADIUS server with incorrect code 122 (invalid request) wrongly, before sending CoA-ACK packet in response to the original CoA-Request that was being processed. In this case, the router should ignore all RADIUS CoA-Request retries and respond only to the original CoA-Request packet. PR1198691
- On MX Series routers with subscriber management feature enabled, after GRES switchover **show network-access aaa statistics radius** command display only zeros and **clear network-access aaa statistics radius** does not clear statistics as it should. It is 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

See Also

- New and Changed Features on page 19
- Changes in Behavior and Syntax on page 35
- Known Behavior on page 38
- Resolved Issues on page 42
- Documentation Updates on page 58
- Migration, Upgrade, and Downgrade Instructions on page 59
- Product Compatibility on page 68

Resolved Issues

This section lists the issues fixed in the Junos OS main 16.2R1 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: 16.2R1 on page 43

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: http://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
- As a workaround, the certificate can be input per rfc7468 section 5.1. Appending the two characters '\' and 'n' to represent newline after the header, base64 data lines, and footer. Or using a cert file with actual newline chars via 'load-key-file'. 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 19
- Changes in Behavior and Syntax on page 35
- Known Behavior on page 38
- Known Issues on page 38
- Documentation Updates on page 58
- Migration, Upgrade, and Downgrade Instructions on page 59
- Product Compatibility on page 68

Documentation Updates

There are no errata or changes in Junos OS Release 16.2R1 documentation for MX Series and T Series.

See Also

- New and Changed Features on page 19
- Changes in Behavior and Syntax on page 35

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 6.1. This feature includes a simplified package naming system that drops the domestic and world-wide naming convention. However, in some of the routers, FreeBSD 6.1 remains the underlying OS for Junos OS. For more details about FreeBSD 10.x, see Understanding Junos OS with Upgraded FreeBSD.

NOTE: **In Junos OS Release 15.1, Junos OS (FreeBSD 10.x) is not available to customers in Belarus, Kazakhstan, and Russia. Customers in these countries need to use the existing Junos OS (FreeBSD 6.1).**

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

Platform	FreeBSD 6.1-based Junos OS	FreeBSD 10.x-based Junos OS
MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 16.2

When upgrading or downgrading Junos OS, always use the **jinstall** package. Use other packages (such as the **jbundle** package) only when so instructed by a Juniper Networks support representative. For information about the contents of the **jinstall** package and details of the installation process, see the Installation and Upgrade Guide and Upgrading Junos OS with Upgraded FreeBSD.

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

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

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

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

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

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

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

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

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

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-32-16.2R1.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot
source/junos-install-mx-x86-64-16.2R1.9-signed.tgz
```

Replace ***source*** with one of the following values:

- ***/pathname***— For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname/pathname***
  - **http://*hostname/pathname***
  - **scp://*hostname/pathname*** (available only for Canada and U.S. version)

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 10.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 10.x, and Junos OS (FreeBSD 6.1) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

Use the **reboot** command to reboot the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: 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 VM Host Installation topic in the *Software Installation and Upgrade Guide*.

NOTE: After you install a Junos OS Release 16.2 **jinstall** package, you cannot issue the **request system software rollback** command to return to the previously installed Junos OS (FreeBSD 6.1) software. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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 *Software Installation and Upgrade Guide*.

Upgrading from Junos OS (FreeBSD 6.1) to Junos OS (FreeBSD 6.1)

Products impacted: MX80, and MX104.

NOTE: Customers in Belarus, Kazakhstan, and Russia must use the following procedure for all MX Series routers running Junos OS Release 16.2.

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

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

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

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

  ```
  user@host> request system software add validate reboot
  source/jinstall-16.2R1.9-domestic-signed.tgz
  ```

- All other customers, use the following command:

  ```
  user@host> request system software add validate reboot
  source/jinstall-16.2R1.9-export-signed.tgz
  ```

Replace ***source*** with one of the following values:

- ***/pathname***— For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp:*//hostname/pathname***
  - **http:*//hostname/pathname***
  - **scp:*//hostname/pathname*** (available only for Canada and U.S. version)

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Use the **reboot** command to reboot the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

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

Upgrade and Downgrade Support Policy for Junos OS Releases

Support for upgrades and downgrades that span more than three Junos OS releases at a time is not provided, except for releases that are designated as Extended End-of-Life (EEOL) releases. EEOL releases provide direct upgrade and downgrade paths— you can

upgrade directly from one EEOL release to the next EEOL release even though EEOL releases generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 14.1, 14.2, 15.1 and 16.1 are EEOL releases. You can upgrade from Junos OS Release 14.1 to Release 15.1 or 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 http://www.juniper.net/support/eol/junos.html.

Upgrading a Router with Redundant Routing Engines

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

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

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

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

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

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

as provider edge (PE) routers. Doing so preserves multicast VPN connectivity throughout the upgrade process.

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

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

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

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

   

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

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

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

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

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

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

   

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

For more information about configuring the draft-rosen Multicast VPN feature, see the *Multicast Protocols Feature Guide for Routing Devices*.

Upgrading Using Unified ISSU

CAUTION:

When In-Service Software Upgrade (ISSU) is performed on MX Series routers from Junos OS 14.2, 15.1, or 16.1 to Junos OS 16.2, the following MPCs and its MICs of both Ethernet and Non-Ethernet variations fail to come online after the ISSU process is complete:

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

Unified in-service software upgrade (ISSU) enables you to upgrade between two different Junos OS releases with no disruption on the control plane and with minimal disruption of traffic. Unified ISSU is only supported by dual Routing Engine platforms. In addition, graceful Routing Engine switchover (GRES) and nonstop active routing (NSR) must be enabled. For additional information about using unified ISSU, see the *High Availability Feature Guide for Routing Devices*.

For information on ISSU support across platforms and Junos OS releases, see the In-Service Software Upgrade (ISSU) Web application.

Downgrading from Release 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 example, if your routing platform is running Junos OS Release 11.4, you can downgrade the software to Release 10.4 directly, but not to Release 10.3 or earlier; as a workaround, you can first downgrade to Release 10.4 and then downgrade to Release 10.3.

For more information, see the *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 19
- Changes in Behavior and Syntax on page 35
- Known Behavior on page 38
- Known Issues on page 38
- Resolved Issues on page 42
- Documentation Updates on page 58
- Product Compatibility on page 68

Product Compatibility

- Hardware Compatibility on page 68

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: http://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 19
- Changes in Behavior and Syntax on page 35
- Known Behavior on page 38

Junos OS Release Notes for PTX Series Packet Transport Routers

These release notes accompany Junos OS Release 16.2R1 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

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

New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 16.2R1 for the PTX Series.

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.

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.

Routing Protocols

- **IS-IS purge originator identification TLV (PTX Series)**—Beginning with Release 15.1F4 and 16.2R1, Junos OS supports RFC 6232, *Purge Originator Identification TLV for IS-IS*, which defines a type, length, and value (TLV) for identifying the origin of a purge initiated by the IS-IS protocol. You can configure this feature to add this TLV to a purge, along with the system ID of the Intermediate System (IS) that has initiated this purge. This makes it easier to locate the origin of the purge and its cause. A new show command **show isis purge log** is introduced to view the purge history and to identify the purge originator.

  [See IS-IS Purge Originator Identification Overview.]

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 73
- Known Behavior on page 75
- Known Issues on page 76
- Resolved Issues on page 76
- Documentation Updates on page 80
- Migration, Upgrade, and Downgrade Instructions on page 80
- Product Compatibility on page 84

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 16.2R1 for the PTX Series.

- Management on page 74
- Network Management and Monitoring on page 74
- Platform and Infrastructure on page 74
- System Logging on page 74

Management

- **Support for status deprecated statement in YANG modules (PTX Series)**—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.

Network Management and Monitoring

- **Possible change is in the object identifier (PTX 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.

Platform and Infrastructure

- **Improvements to MIB validation during Junos build (libjsnmp) (PTX Series)**—There are many warnings which can be ignored if produced while MIBs are compiling. Following are some warnings that you need to consider as 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 in jnx-chas-defines.

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

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

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.

See Also

- New and Changed Features on page 70
- Known Behavior on page 75
- Known Issues on page 76
- Resolved Issues on page 76
- Documentation Updates on page 80
- Migration, Upgrade, and Downgrade Instructions on page 80
- Product Compatibility on page 84

Known Behavior

There are no known limitations in Junos OS Release 16.2R1 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 70
- Changes in Behavior and Syntax on page 73
- Known Issues on page 76
- Resolved Issues on page 76
- Documentation Updates on page 80
- Migration, Upgrade, and Downgrade Instructions on page 80
- Product Compatibility on page 84

Known Issues

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

- Major errors might be seen on MPC3/FPC3 with 1X100 and 5x100 DWDM MIC/PIC. root@abc-r1-re0> show chassis alarms no-forwarding 1 alarms currently active Alarm time Class Description 2016-01-1 11:22:44 UTC 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

See Also

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.

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 70

Documentation Updates

There are no errata or changes in Junos OS Release 16.2R1 documentation for PTX Series.

See Also

- New and Changed Features on page 70
- Changes in Behavior and Syntax on page 73
- Known Behavior on page 75
- Known Issues on page 76
- Resolved Issues on page 76
- Migration, Upgrade, and Downgrade Instructions on page 80
- Product Compatibility on page 84

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

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 for Routing Devices.

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

The download and installation process for Junos OS Release 16.2R1 is different from previous Junos OS releases.

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

    

    NOTE: After you install a Junos OS Release 16.2R1 jinstall package, you cannot issue the request system software rollback command to return to the previously installed software. Instead you must issue the request system software add validate command and specify the jinstall package that corresponds to the previously installed software.

    The validate option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is

a different release. Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes. Rebooting occurs only if the upgrade is successful.

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

```
user@host> request system software add validate reboot source/jinstall-16.2
B.1-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot source/jinstall-16.2
R.1-export-signed.tgz
```

Replace the ***source*** with one of the following values:

- ***/pathname***— For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp:*//hostname/pathname***
  - **http:*//hostname/pathname***
  - **scp:*//hostname/pathname*** (available only for Canada and U.S. version)

The **validate** option validates the software package against the current configuration as a prerequisite to adding the software package to ensure that the router reboots successfully. This is the default behavior when the software package being added is a different release.

Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process can take 5 to 10 minutes.

Rebooting occurs only if the upgrade is successful.

NOTE: 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 *Software Installation and Upgrade Guide*.

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

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 70
- Changes in Behavior and Syntax on page 73
- Known Behavior on page 75
- Known Issues on page 76
- Resolved Issues on page 76
- Documentation Updates on page 80
- Product Compatibility on page 84

Product Compatibility

- Hardware Compatibility on page 85

Hardware Compatibility

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

To determine the features supported on PTX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: http://pathfinder.juniper.net/feature-explorer/.

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 70
- Changes in Behavior and Syntax on page 73
- Known Behavior on page 75
- Known Issues on page 76
- Resolved Issues on page 76
- Documentation Updates on page 80
- Migration, Upgrade, and Downgrade Instructions on page 80

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: http://prsearch.juniper.net.

Juniper Networks Feature Explorer is a Web-based application that helps you to explore and compare Junos OS feature information to find the correct software release and hardware platform for your network. Find Feature Explorer at: http://pathfinder.juniper.net/feature-explorer/.

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

Documentation Feedback

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

- Online feedback rating system— On any page at the Juniper Networks Technical Documentation site at http://www.juniper.net/techpubs/index.html, simply click the stars to rate the content, and use the pop-up form to provide us with information about your experience. Alternately, you can use the online feedback form at http://www.juniper.net/techpubs/feedback/.
- E-mail— Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

Technical product support is available through the Juniper Networks Technical Assistance Center (JTAC). If you are a customer with an active J-Care or JNASC support contract, or are covered under warranty, and need postsales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies— For a complete understanding of our JTAC procedures and policies, review the JTAC User Guide located at http://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties— For product warranty information, visit http://www.juniper.net/support/warranty/.
- JTAC Hours of Operation — The JTAC centers have resources available 24 hours a day, 7 days a week, 365 days a year.

Self-Help Online Tools and Resources

For quick and easy problem resolution, Juniper Networks has designed an online self-service portal called the Customer Support Center (CSC) that provides you with the following features:

- Find CSC offerings: http://www.juniper.net/customers/support/
- Search for known bugs: http://www2.juniper.net/kb/
- Find product documentation: http://www.juniper.net/techpubs/
- Find solutions and answer questions using our Knowledge Base: http://kb.juniper.net/
- Download the latest versions of software and review release notes: http://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: http://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: http://www.juniper.net/company/communities/
- Open a case online in the CSC Case Management tool: http://www.juniper.net/cm/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) tool located at https://tools.juniper.net/SerialNumberEntitlementSearch/.

Opening a Case with JTAC

You can open a case with JTAC on the Web or by telephone.

- Use the Case Management tool in the CSC at http://www.juniper.net/cm/.
- Call 1-888-314-JTAC (1-888-314-5822 toll-free in the USA, Canada, and Mexico).

For international or direct-dial options in countries without toll-free numbers, visit us at http://www.juniper.net/support/requesting-support.html.

If you are reporting a hardware or software problem, issue the following command from the CLI before contacting support:

```
user@host> request support information | save filename
```

To provide a core file to Juniper Networks for analysis, compress the file with the **gzip** utility, rename the file to include your company name, and copy it to **ftp.juniper.net/pub/incoming**. Then send the filename, along with software version information (the output of the **show version** command) and the configuration, to **support@juniper.net**. For documentation issues, fill out the bug report form located at https://www.juniper.net/cgi-bin/docbugreport/.

Revision History

14 February 2019—Revision 8, Junos OS Release 16.2R1– ACX Series, EX Series, MX Series, PTX Series, and T Series.

2 August 2018—Revision 7, Junos OS Release 16.2R1– 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, and T Series.

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

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

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

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

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