Release Notes: Junos® OS Release 18.3R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

26 August 2021

Contents

Introduction

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

These release notes accompany Junos OS Release 18.3R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the ACX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

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

Release 18.3R2 New and Changed Features

- There are no new features or enhancements to existing features for ACX Series Universal Metro Routers in Junos OS Release 18.3R2.

Release 18.3R1-S1 New and Changed Features

IN THIS SECTION

Class of Service (CoS)

- **Support for deep buffer and drop profiles (ACX5448)**—Starting with Junos OS Release 18.3R1-S1, ACX5448 router supports the use of external DRAM memory, along with its on-chip memory, for scheduling and queuing different virtual output queues (VOQs). ACX5448 router also supports different WRED profiles for configuring drop profiles for queues.

  NOTE: ACX5448 router does not support buffering for IRB multicast traffic and therefore CLIs for configuring multicast is not supported.

  [See Configuring Shared and Dedicated Buffer Memory Pools.]

Timing and Synchronization

- **Support for PTP boundary clocks for phase and time synchronization (ACX5448)**—Starting with Junos OS Release 18.3R1-S1, ACX5448 router supports PTP boundary clocks for phase and time synchronization using IEEE-1588 Precision Timing Protocol (PTP). This feature also supports:
  - PTP over IPv4 (IEEE-1588v2)
  - PTP ordinary and boundary clocks
  - One step clock mode operation for PTP Master
  - 10Mhz and 1PPS output for measurement purpose

  NOTE: All PTP packets uses the best-effort queue instead of network control queue.

  The ACX5448 router does not support the following:

  - Hybrid mode
  - Boundary clock performance complying to G.8273.2
  - Dual tagged PTP over IPv4

  [See IEEE 1588v2 PTP Boundary Clock Overview.]

Release 18.3R1 New and Changed Features

IN THIS SECTION

Hardware

- **New fixed-configuration universal metro router (ACX Series)**—Starting in Junos OS Release 18.3R1, the ACX6360 is a new fixed-configuration router that provides full IP/MPLS stack and secure packet optical transport convergence. It features a compact, 1U form factor that can perform as either a transponder or a router. It can supply either muxponder-like pass-through connection of client interface traffic to line optical interfaces or IP/MPLS routing services. The ACX6360 has 20 QSFP28 ports and 8 CFP2 ports. When the ACX6360 is configured as a router, the 20 QSFP28 ports can be configured as 10 Gbps, 40 Gbps, or 100 Gbps. When the ACX6360 is configured as a transponder, the 20 QSFP28 ports can be configured as 100 Gbps. The 8 CFP2 ports can be configured as 100 Gbps or 200 Gbps.

  [See ACX6360 Documentation.]

Authentication, Authorization and Accounting

- **Support for password change policy enhancement (ACX Series)**—Starting in Junos OS Release 18.3R1, the Junos password change policy for local user accounts is enhanced to comply with certain additional password policies. As part of the policy improvement, you can configure the following:
  - **minimum-character-changes**—The number of characters by which the new password should be different from the existing password.
  - **minimum-reuse**—The number of older passwords, which should not match the new password.

[See password.]

Interfaces and Chassis

- **Support for pre-FEC BER monitoring (ACX6360)**—Starting in Junos OS Release 18.3R1, you can monitor the condition of an OTN link on an ACX6360 router by using the pre-forward error correction (pre-FEC) bit error rate (BER). The ACX6360 router uses FEC to correct bit errors in the received data. As long as the pre-FEC BER is below the FEC limit, all bit errors are successfully identified and corrected and, therefore, no packet loss occurs. The router monitors the pre-FEC BER on each port, which provides an early indication of possible link degradation. By configuring an appropriate pre-FEC BER threshold and interval, you enable the ACX6360 router to take preemptive action before the FEC limit is reached.

  [See Understanding Pre-FEC BER Monitoring and BER Thresholds.]

- **ACX6360 routers support router mode and transponder mode**—Starting in Junos OS Release 18.3R1, ACX6360 routers support two modes - optical router mode (chassis model: ACX6360-OR) and optical transponder mode (chassis model: ACX6360-OX). While the ACX6360 in optical router mode supports routing centric features, in optical transponder mode, the device functions as an optical transponder, which does not support the routing features. In optical transponder mode, up to 16 cross-connects are created between QSFP28 client ports and CFP2-DCO ports by default. To enable optical transponder mode, install the ACX-OX version of the Junos OS VM host image (ACX-OX) in the chassis. To enable optical router mode, install the ACX-OR version of the Junos OS VM host image (ACX-OR). You can use the Junos OS CLI command **request vmhost software add** to install the Junos VM host images.

  [See Understanding Router Mode and Transponder Mode on ACX6360.]

Junos OS XML API and Scripting

- **Support for Python language for commit, event, op, and SNMP scripts (ACX5048 and ACX5096)**—Starting in Junos OS Release 18.3R1, you can author commit, event, op, and SNMP scripts in Python on devices that include the Python extensions package in the software image. Creating automation scripts in Python enables you to take advantage of Python features and libraries as well as leverage Junos PyEZ APIs to perform operational and configuration tasks on devices running Junos OS. To enable execution of Python automation scripts, which must be owned by either root or a user in the Junos OS **super-user** login class, configure the **language python** statement at the **[edit system scripts]** hierarchy level, and configure the filename for the Python script under the hierarchy level appropriate to that script type. Supported Python versions include Python 2.7.

  [See Understanding Python Automation Scripts for Devices Running Junos OS.]

Junos Telemetry Interface

- **Support for the Junos Telemetry Interface (ACX6360, MX Series, and PTX Series)**—Starting with Junos OS Release 18.3R1, Junos Telemetry Interface support is available for the ACX6360 Universal Metro Router and MX Series and PTX Series routers with a CFP2-DCO optics module that provides a high-density, long-haul optical transport network (OTN) transport solution with MAC capability.

  You can provision sensors to export telemetry data to an outside collector.

The following native (UDP) and gRPC sensors can be provisioned for ET (100-Gigabit Ethernet) interfaces and OT interfaces:

- /junos/system/linecard/optical
- /junos/system/linecard/otn

To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

[See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **Expanded physical interface queue and traffic statistics sensors for Junos Telemetry Interface (JTI) (ACX Series)**—Starting with Junos OS Release 18.3R1, additional resource paths are added to stream physical (IFD) statistics.

  Prior to Junos OS Release 18.3R1, both traffic and queue statistics for physical interfaces (IFD) are sent out together using the resource path **/interfaces** for gRPC streaming (which is internally used to create **/junos/system/linecard/interface/**) or **/junos/system/linecard/interface/** for UDP (native) sensors.

  Now, traffic and queue statistics can be delivered separately. Doing so can reduce the reap time for non-queue data for platforms supporting Virtual Output Queues (VOQ).

  The following UDP resource paths can be configured:

  - **/junos/system/linecard/interface/** is the existing resource path (no change). Traffic and queue statistics are sent together.
  - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
  - **/junos/system/linecard/interface/queue/** exports queue statistics.

  The gRPC resource path **/interfaces** now has the following behavior:

  - In releases prior to Junos OS 18.3R1, it delivers all IFD traffic and queue statistics. In Junos OS 18.3R1 and higher, it delivers statistics in two sensors:
    - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
    - **/junos/system/linecard/interface/queue/** exports queue statistics.

  To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

  [See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

Layer 2 Features

- **Support for Layer 2 RFC2544 reflection (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports the Layer 2 RFC2544 reflector functionality to reflect the test packets back to the network. This feature is useful for verifying the connectivity and fault isolation. It can be used for performance measurement where the RFC2544 feature functionality can loopback the packets to a measuring device. The RFC2544 feature functionality supports:
  - RFC2544 egress Layer 2 reflection functionality for family bridge.
  - Multiple RFC2544 reflection sessions.
  - Reflection on 1G/10G/40G/Ch10G/Ch25G/100G ports.
  - Ethernet Layer 2 frames to carry IP/UDP packets for RFC2544 reflection.

  ACX5448 router do not support the following RFC2544 features:
  - Any interface in the bridge domain matching the bridge VLAN identifier is not supported.
  - Multiple simultaneous sessions with multiple VLAN bridges are not supported.
  - Multiple test sessions cannot exceed 100G bandwidth.
  - IPv6 reflection.
  - IPV6 filter support to identify the loopback stream.
  - RFC 2544 reflection functionality for family **ccc** (PWE reflection) and family **inet** (Layer 3 IPv4 reflection).
  - Reflection without MAC swap and MAC overwrite is not supported.
  - Reflection on ELINE/ELAN services.

[See RFC 2544-Based Benchmarking Tests Overview.]

MPLS

- **Support for MPLS fast reroute and unicast reverse path forwarding (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports MPLS fast reroute (FRR) and unicast reverse-path forwarding (uRPF). Fast reroute provides redundancy for an LSP path. When you enable fast reroute, detours are precomputed and preestablished along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Fast reroute protects traffic against any single point of failure between the ingress and egress routers (or switches).

  A unicast reverse-path-forwarding (RPF) check is a tool to reduce forwarding of IP packets that might be spoofing an address. A unicast RPF check performs a route table lookup on an IP packet's source address, and checks the incoming interface. The router determines whether the packet is arriving from a path that the sender would use to reach the destination. If the packet is from a valid path, the router forwards the packet to the destination address. If it is not from a valid path, the router discards the packet. Unicast RPF is supported for the IPv4 and IPv6 protocol families, as well as for the virtual private network (VPN) address family.

  [See Fast Reroute Overview and Guidelines for Configuring Unicast RPF on ACX Series Routers.]

Multicast

- **Support for IPv6 multicast using Multicast Listener Discovery protocol (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports IPv6 multicast using Multicast Listener Discovery (MLD) protocol. To support multicast data delivery, ACX5448 router supports 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.

  [See Understanding MLD, IPv6 Multicast Flow, and Enabling MLD.]

OAM

- **Support for Operations, Administration, and Management (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports the following Operations, Administration, and Management (OAM) feature standards:

  - IEEE Standard 802.1ag, also known as connectivity fault management (CFM).
  - ITU-T Recommendation Y.1731, which uses different terminology than IEEE 802.1ag and defines Ethernet service OAM features for fault monitoring, diagnostics, and performance monitoring.
  - IEEE Standard 802.3ah for OAM link fault management (LFM).

  The OAM feature in ACX5448 router includes support for maintenance endpoints (MEPs). MEPs can be up MEPs or down MEPs. A MEP can be configured to support continuity check message (CCM), loopback message, delay measurement, and synthetic loss message (SLM) message types. ACX5448 router also supports OAM for VPLS.

> **NOTE:** ACX5448 router do not support maintenance association intermediate point (MIP).

[See Ethernet OAM Connectivity Fault Management and Understanding Ethernet OAM Link Fault Management for ACX Series Routers.]

Routing Policy and Firewall Filters

- **Support for firewall filters and policers (ACX5448)**—Starting with Junos OS Release 18.3R1, you can configure firewall filters on packets (families such as bridge domain, IPv4, IPv6, CCC, MPLS, VPLS) based on packet match conditions with the support of external TCAM in ACX5448 router. Along with the match conditions, actions such as count, discard, log, syslog, policer are performed on the packets that match the filter. You can configure policers and attach them to a firewall term. This feature also supports configuring ARP policer, forwarding table filters, and policy-based routing.

  This feature enables scaling the family filters of the firewall functionality in the ingress direction.

  The following ingress family filters can be scaled based on the availability of external-tcam:

  - family **ethernet-switching**
  - family **ccc**
  - family **inet**
  - family **inet6**
  - family **mpls**
  - family **vpls**

  The loopback (**Lo0**) filters, family **any**, and other module applications continue to use internal-tcam and can reach maximum of the internal-tcam.

  [See Firewall Filter Match Conditions and Actions on ACX Series Routers Overview.]

Routing Protocols

- **Support for Virtual Router Redundancy Protocol (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports Virtual Router Redundancy Protocol (VRRP) as per RFC 3798 VRRP version 2 and RFC 5798 VRRP version 3. ACX5448 router also supports configuring VRRP over aggregated Ethernet and integrated routing and bridging (IRB) interfaces.

  The following limitations apply while configuring VRRP on ACX5448 router:

  - Configure a maximum of 16 VRRP groups.
  - Interworking of VRRP version 2 and VRRP version 3 is not supported.
  - VRRP delegate processing is not supported.
  - VRRP version 2 authentication is not supported.

[See Understanding VRRP.]

Timing and Synchronization

- **Support for frequency synchronization using synchronous Ethernet protocol (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports frequency synchronization using synchronous Ethernet (SyncE) protocol, with ESMC support as per the ITU-T standard G.8262/G.8264. This feature also supports 10Mhz and PPS output for measurement purpose.

  [See Clock Sources for ACX Series.]

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (ACX Series)**—Starting in Junos OS Release 18.3R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database. The ephemeral database provides a fast programmatic interface that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. The device's active configuration is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database. Ephemeral configuration data is volatile and is deleted upon rebooting the device.

  [See Understanding the Ephemeral Configuration Database.]

VPN

- **Support for Layer 3 VPN and IPv6 VPN Provider Edge Router (6VPE) over MPLS (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports Layer 3 VPN and IPv6 VPN provider edge router (6VPE) support over MPLS. Layer 3 VPNs are based on RFC 4364 that defines a mechanism by which service providers can use their IP backbones to provide VPN services to their customers. ACX5448 router, acting as a VPN provider edge router, provides IPv6 forwarding over MPLS. 6VPE adds IPv6 support to the current IPv4 MPLS by transporting IPv6 across MPLS core.

  [See Understanding Layer 3 VPNs.]

VLAN Infrastructure

- **Support for VPLS features (ACX5448)**—Starting with Junos OS Release 18.3R1, ACX5448 router supports full-mesh VPLS domain deployment. ACX5448 router supports interworking of both BGP as well as LDP-based VPLS. BGP can be used only for auto-discovery of the VPLS PEs, while LDP signaling for VPLS connectivity.

  The following VPLS configurations are supported:

  - VPLS domains
  - VLAN identifier and VLAN maps
  - MAC learning
  - Logical interface support

- Control protocol support
- Interworking of LDP and BGP VPLS
- VCCV BFD support
- Firewalls and filter support

[See Introduction to VPLS.]

SEE ALSO

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for the ACX Series routers.

Junos OS XML, API, and Scripting

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (ACX Series)**—Starting in Junos OS Release 18.3R1, Junos OS does not support configuring an MD5 or SHA-1 checksum hash to verify the integrity of local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) scripts or support using an MD5 or SHA-1 checksum hash with the **op url *url* key** option to verify the integrity of remote op scripts.

Network Management and Monitoring

- **Junos OS does not support management of YANG packages in configuration mode (ACX Series)**—Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages using the **run** command in configuration mode is not supported.
- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (ACX Series)**—Starting in Junos OS Release 18.3R2, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

Platform and Infrastructure

- **DMA recovery mechanism (ACX Series)**—A recovery mechanism has been introduced that is triggered in case the router enters an Idle state on any DMA channels. The recovery mechanism reboots the PFE to recover from Idle state.

  The following recovery message is logged in the RE syslog message:

  ```
  CHASSISD_FPC_ASIC_ERROR: <FPC 0> ASIC Error detected errorno 0x0000ffff FPC
  restart initiated
  ```

  The following recovery message is logged in the PFE syslog message:

  ```
  BCM DMA channel error detected
  Resetting the PFE
  ```

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (ACX Series)**—Starting in Junos OS Release 18.3R1, the jdhcpd process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
  - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
  - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

  The new behavior is as follows:

  1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
  2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

  In earlier releases:

  1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
  2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

  [See Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview.]

SEE ALSO

New and Changed Features | 12

Known Behavior | 24

Known Issues | 25

Resolved Issues | 29

Documentation Updates | 32

Migration, Upgrade, and Downgrade Instructions | 33

Product Compatibility | 34

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 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.

General Routing

- Upon classifying the Layer 3 packets, DSCP might not be preserved or lost at the egress due to the limitations of the forwarding asic. PR1322142
- The typical transponder propagates the pause frame received from client port to line port. For this Packet Forwarding Engine transponder, this functionality is not supported. PR1371603
- No new commit errors will be thrown when **buffer-size temporal** is configured along with **shared-buffer maximum**. PR1371828
- Telemetery infrastructure does not support interface filtering capability. Therefore, once you enable a particular sensor for telemetry, it is turned-on for all the interfaces. PR1371996
- For -et interfaces, only PRE_FEC_SD defect will be raised no OTN alarm will be raised. PR1371997
- If you configure an invalid sandbox configuration, CCC functionality will break after reboot or upgrade. Sandbox configuration is always done initially by default and you must not modify this configuration. PR1373375
- L2 rewrite on outgoing MPLS packets is not supported. PR1376001
- When the system is commissioned first time after upgrade, root authentication configuration needs to be entered. All the default cross-connect configurations done by the script is not saved in configuration till system root authentication configuration is entered. This is a Junos OS product feature. So, if user displays the cross-connect configuration before configuring root authentication then cross-connect configuration would not be visible. Current product limitations are: 1. System root authentication configuration is needed after system is commissioned prior to the init script run otherwise the cross connect installation might fail. 2. If the existing CCC configurations (user defined cross connects) are different than the defaults, the configurations might be lost and will be replaced by default cross-connects after the software upgrade. 3. Software upgrade needs no-validate option during installation. PR1376780

- The **static-cak** encryption does not work between two ACX-OX transponder nodes. PR1389802
- For the ACX6360 TIC we only have 8 CFP2-DCO ports so chassis beacon show/requests to ports larger than 7 will not work (as the ports do not exist) but will also not report an error. **user@host> request chassis beacon fpc 0 pic-slot 1 port 15 on FPC 0 PIC 1 PORT 15 ON user@host> show chassis beacon fpc 0 pic-slot 1 port-range lower-limit 0 upper-limit 15 FPC 0 PIC 1 PORT 0 ON FPC 0 PIC 1 PORT 1 ON FPC 0 PIC 1 PORT 2 ON FPC 0 PIC 1 PORT 3 ON FPC 0 PIC 1 PORT 4 ON FPC 0 PIC 1 PORT 5 ON FPC 0 PIC 1 PORT 6 ON FPC 0 PIC 1 PORT 7 ON FPC 0 PIC 1 PORT 8 ON FPC 0 PIC 1 PORT 9 ON FPC 0 PIC 1 PORT 10 OFF FPC 0 PIC 1 PORT 11 OFF FPC 0 PIC 1 PORT 12 OFF FPC 0 PIC 1 PORT 13 OFF FPC 0 PIC 1 PORT 14 OFF FPC 0 PIC 1 PORT 15 ON**. PR1399335
- The policers applied in IRB will work appropriately when the member links of an aggregated Ethernet interface is in the same core file applied in the aggregated Ethernet interface. The physical interface might generate a core file mapping: **xe-0/0/0 -to- xe-0/0/23 -> CORE 0 xe-0/0/24 -to- xe-0/0/47 -> CORE 1 et-0/1/0 -> CORE 1 et-0/1/1 -> CORE 1 et-0/1/2 -> CORE 0 et-0/1/3 -> CORE 0**. The policers applied in IRB will work appropriately when the member links of a bridge domain (BD) is in the same core file. PR1403315
- If user configures an invalid speed configuration on TIC ports (PIC slot 1) on ACX6360-OR or ACX6360-OX, the TIC interfaces are not created. PR1403546

SEE ALSO

New and Changed Features | **12**

Changes in Behavior and Syntax | **21**

Known Issues | **25**

Resolved Issues | **29**

Documentation Updates | **32**

Migration, Upgrade, and Downgrade Instructions | **33**

Product Compatibility | **34**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 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.

General Routing

- When the ACX2100 and ACX2200 are used as ingress PE routers for L2 circuit connections, and the PE-CE interface (UNI) is an aggregated Ethernet interface, then upon MPLS path switchover, the traffic might be silently dropped or discarded. PR1194551
- Enhancement of logical interface scale beyond 1000 on ACX5000 platform is not available in mainline release starting from Junos OS Release 17.1 and later. PR1229492
- On ACX5448 routers, when 1-Gigabit SFP is plugged in the router, autonegotiation is enabled by default. There is no functional impact. Only the CLI **show interfaces <intf-name> extensive** command output shows the autonegotiation field as disabled. PR1343679
- There is a conflict when an LACP packet come in an untagged/prio-tagged VPLS logical interface. In the earlier stage of the pipeline, filter entry to snoop an LACP packet takes higher precedence over filter entry to assign SVP/SrcG port for the untagged/prio-tagged VPLS logical interface. Since the "interface-specific/input-list" firewall matches SVP/SrcGport in the later stage of the pipeline, the LACP packets are not hitting the firewall. PR1346380
- The logical interface classifier information should not be shown in the output of the **show class-of-service interface <ifd>** on the ACX5000 line. PR1353828
- On an ACX5448 chassis with loss priority configured as medium-low or medium-high, the rewrite rule gets applied for loss priority low. PR1358721
- Remote fault signalling is not supported for 1-Gigabit fiber SFP during autonegotiation. The following cosmetic log errors are seen for **show interfaces extensive** command. **Link partner: Link mode: Full-duplex, Flow control: None, Remote fault: Down, Reason: Link partner offline. RFI ignored since AN is in default mode**. PR1362490
- Dedicated minimum buffers are reserved for some queues according to the Junos OS working model. These buffers are always available to those queues irrespective of the traffic pattern throughout the system. When the **clearing stat** statement is used, these values are visible. This cosmetic or minor issue has no functional impact. PR1367978

- Because of a race condition, in which the **class-of-service** configuration request for an interface is received before the e1-interface is created, a circuit with specified class-of-service parameters is created. Because of this, the interface creation fails, resulting in traffic not flowing on the e1-interface and then (if e1-interfaces are further disabled or enabled) a core file is generated. PR1378747
- The dedicated buffer for bytes/packets sometimes exceeds the maximum threshold value under the **show class-of-service packet-buffer usage** command output. As per DNX architecture, reserved buffer is not limited to the OCB buffer limit (16 MB), so whenever the buffer goes beyond 16 MB, DNX punts the packet to the DRAM instead of dropping it. This is as per design or behavior. PR1379713
- Host bound traffic might be affected and lt interface might go down in ACX Series routers. PR1382166
- When packets are sent from Layer 2 to Layer 2 and when you apply MF classifier, all packets are put into the correct queue on the egress interface but they are dropped. As a workaround, avoid the **loss-priority high** action in the firewall filters (MF classifiers). PR1388731
- On the ACX5000 line, in Junos OS Release 17.3 and later releases, the Packet Forwarding Engine syslog frequently shows the following error message: **acx_cos_tcp_bind_queues:736 parent acx_cos_tcp_ifd for ifd:ae0 doesn't exist for ifl:549**. In Junos OS Release 17.3R3-S1, the error logs appear only from time to time, and this can be related with to an interface flap. In Junos OS Release 18.1R3, the logs appear constantly, without any interface flap. PR1392088
- Explicit swap-push map operations are now introduced on VPLS logical interfaces in ACX5000. This is already supported as part of implicit map operations or routing instance-level configurations. PR1398118
- A jnxIfOtnOperState trap notification is sent for all ot-interfaces. This is a day-1 issue. PR1406758
- Policer discarded packets are marked in black color (black is color-internal to hardware pipeline). Black color is used to discard the packets in the pipeline. These packets are not really enqueued into the queues (VoQs) in hardware. The hardware queue statistics shows the packets as discarded. However, both actual-enqueued and the discarded counts are shown as queue statistics in software. This is a software queue statistics show issue. PR1414887
- Packets transmitted in a queue are not as expected when testing IEEE-802.1ad inner classifier at the ingress and IEEE-802.1ad rewrite at the egress with various events. PR1422515
- Copying images from WAN interface to Routing Engine of ACX5448 router takes long time. PR1422544

Interfaces and Chassis

- When an unnumbered interface is binding to an interface that has more than one IP address and one of the IP addresses is deleted, the family inet of the unnumbered interface might be deleted. The issue results in traffic loss for all the services that rely on the family inet of the unnumbered interface. Configuring **preferred-source-address** on the unnumbered interface will prevent deletion of the IP address thereby avoiding the deletion of the family inet of the unnumbered interface. PR1412534

Layer 2 Features

- On Junos OS ACX5000, on the interfaces where LLDP is disabled (commit) and there is a change on an interface in the next commit, the l2cpd sends the message to disable LLDP on all the interfaces to the kernel. The kernel then tries to remove the implicit filters, which return ENOENT, since the entries are disabled during the first commit. PR1400606

MPLS

- Packets transmitted in a queue are not as expected when testing IEEE-802.1ad inner classifier at the ingress and IEEE-802.1ad rewrite at the egress with various events. PR1432138

SEE ALSO

New and Changed Features | **12**

Changes in Behavior and Syntax | **21**

Known Behavior | **24**

Resolved Issues | **29**

Documentation Updates | **32**

Migration, Upgrade, and Downgrade Instructions | **33**

Product Compatibility | **34**

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.3R2

Class of Service

- Error message **STUCK_BUFF : port_sp not empty for port 35 sp 1 pkts:1** is seen when a lag bundle is configured with 64 lag links. PR1346452

General Routing

- 10G interface fault detection behavior changed. PR1223457
- ACX Series routers supports dual-tagged and untagged packets Layer 3 traffic. PR1307666
- Port XE-0/3/0 did not come up. PR1328207
- ARP request is getting dropped and not forwarded to the NNI interface queue when the CoS configuration has temporal buffer size. PR1363153
- On ACX5000 routers, the log message **fpc0 (acx_rt_ip_uc_lpm_install:LPM route add failed) Reason : Invalid parameter** is seen after configuring **lpm-profile**. PR1365034
- VPLS with **vlan-id-list** does not work when the link between a PE device and a CE device is an aggregated Ethernet interface with a single member link and child physical interface flap. PR1365894
- **LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** prints while commit on configuration prompt. PR1376665
- The fxpc might crash after an interface is changed on ACX5000 line of routers. PR1378155
- On ACX5448 routers, channelized 25-Gbps et-interfaces might not come up after **chassis-control restart**. PR1379288
- The L2 circuit might stop forwarding traffic when one core interface flaps. PR1381487

- On ACX6360 routers, the timestamp is incorrect for BER statistics after clearing. PR1386253
- The **request chassis beacon** CLI command is not working for pic-slot 1 (that is, CFP2 ports). PR1386711
- On ACX5448 routers, 100-Gbps link FEC is enabled by default on 100-Gbps LR4. PR1389518
- On ACX Series platforms, the **forwarding-option dhcp-relay forward-only** statement stops working and the DHCP packets are dropped. PR1392261
- Certain builds of Junos OS do not allow you to upgrade or commit configuration changes when the SI service interface is used. PR1393729
- On ACX Series routers, the MTU is not properly applied and the output of **ping mpls l2circuit sweep** is giving lower values than expected. PR1393947
- ACX Series routers do not support the **physical-interface-filter** command in egress direction for any filters. It supports the **interface-specific** command only. PR1395362
- On ACX5048 routers, the RPM RFC2544-benchmarking test fails to start. PR1395730
- Error message **ACX_PFE_ERROR: dnx_cfm_bd_endpoint_create: Failed to destroy the remote endpoint, Endpoint id 0x2001001, Entry not found** is logged. PR1397878
- CFM adjacency is not going down with distinct intervals. PR1397883
- Error message **ACX_ASIC_PROGRAMMING_ERROR: dnx_cfm_bd_endpoint_create: Failed to create the local endpoint Invalid parameter** has been logged on peer node. PR1397951
- **Output packet error Count** is incrementing on 40-Gigabit Ethernet and 100-Gigabit Ethernet ports. PR1398270
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. PR1398398
- Dynamic tunnel is not supported on ACX Series routers. PR1398729
- On ACX5448 routers, it is not possible to configure bridge domain (BD) more than 1024, using 100-Gigabit and aggregated Ethernet interface in bridge domain (BD). PR1399214
- FPC might crash after offline/online of MIC-3D-16CHE1-T1-CE-H. PR1402563
- ACX5448 supports TrTCM policer configuration parameters as per RFC4115. PR1405798
- ACX Series routers drop DNS responses that contain an underscore. PR1410062
- The aggregated Ethernet interface TWAMP history statistics verification on client is not getting expected **Request Timed Out** error. PR1411344
- Number of **inet-arp policers** implemented on ACX5000 has been increased from 16 to 64. PR1413807
- Swap memory is not initialized on boot on ACX5048. PR1415898

Services Applications

- The spd might crash when **any-ip** is configured in the from clause of the NAT rule with the static translation type. PR1391928

Resolved Issues: 18.3R1

General Routing

- Several error logs are seen on ACX5048 router when the link in the primary path of LSP is flapped. PR1204714
- A wrong packet statistics is reported in ifHCInUcastPkts OID. PR1306656
- With **auto-installation usb** configured, interface related commits might not take effect due to dcd error. PR1327384
- CoS is wrongly applied on the Packet Forwarding Engine leading to egress traffic drop. PR1329141
- The aggregated Ethernet load balancing based on layer 4 information is not working if ports are in different cores of hardware. PR1332448
- The DHCP negotiations might fail and eventually cause outage if scaling number of DHCP clients reboot at the same time. PR1335957
- The ARP-reply packet might be dropped in a l2-circuit secondary path when using IEEE-802.1 classifier. PR1341126
- NAT might not work and the spd might crash. PR1346546
- On the ACX5448 router, DHCP bindings are not received for both DHCP v4 and v6 as RIO is dropping the DHCP Packets. PR1347906
- The fxpc process might crash on the Packet Forwarding Engine due to the **show pfe context_vlan** command. PR1349721
- The required number of the IGMP SNOOPING Membership reports are not received on ACX5448 router. PR1351422
- DHCP Bindings are not received for both DHCP v4 and v6 as ACX5448 router is dropping the DHCP Packets on aggregated Ethernet interfaces. PR1353887
- On ACX routers, the ARP policer for logical interfaces is not working. PR1356170
- ACX is wrongly allowing to configure higher values in **burst-size-limit** than what the hardware can support. PR1361482
- FEC PM error counters are accumulating instead of resetting after a bin rollover. PR1363270
- An ACK5000 routers, IPsec SA as OSPFv3 authentication is not working in Junos OS Releases 16.2R2 and 17.3R2. PR1363487

- The **show chassis hardware** commands display inconsistent values for PEMs and fans. PR1364224
- The 'commit' or 'commit check' operations might fail due to the **cannot have lsp-cleanup-timer without lsp-provisioning** error. PR1368992

Layer 2 Ethernet Services

- DHCPv6 relay ignores replies from the server when renewing. PR1354212

SEE ALSO

New and Changed Features | **12**

Changes in Behavior and Syntax | **21**

Known Behavior | **24**

Known Issues | **25**

Documentation Updates | **32**

Migration, Upgrade, and Downgrade Instructions | **33**

Product Compatibility | **34**

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 for the ACX Series documentation.

SEE ALSO

New and Changed Features | **12**

Changes in Behavior and Syntax | **21**

Known Behavior | **24**

Known Issues | **25**

Resolved Issues | **29**

Migration, Upgrade, and Downgrade Instructions | **33**

Changes in Behavior and Syntax | **245**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for the ACX Series Router. Upgrading or downgrading Junos OS might take several minutes, 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

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 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

For information about software installation and upgrade, see the Installation and Upgrade Guide.

SEE ALSO

- New and Changed Features | 12
- Changes in Behavior and Syntax | 21
- Known Behavior | 24

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the EX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

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

NOTE: The following EX Series switches are supported in Release 18.3R2: EX2300, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

Release 18.3R2 New and Changed Features

- There are no new features or enhancements to existing features for EX Series switches in Junos OS Release 18.3R2.

Release 18.3R1 New and Changed Features

IN THIS SECTION

Hardware

- **EX4650-48Y switches**—Starting with Junos OS Release 18.3R1, the EX4650-48Y switch is available as a fixed-configuration switch with the following built-in ports:
  - Forty-eight 25-Gigabit Ethernet ports that can operate at 1-Gbps, 10-Gbps, or 25-Gbps speed and support SFP, SFP+, or QSFP28 transceivers.
  - Eight 100-Gigabit Ethernet ports that can operate at 40-Gbps or 100-Gbps speed and support QSFP+ or QSFP28 transceivers. When these ports operate at 40-Gbps speed, you can configure four 10-Gbps interfaces and connect breakout cables, increasing the total number of supported 10-Gbps ports to 80. When these ports operate at 100-Gbps speed, you can configure four 25-Gbps interfaces and connect breakout cables, increasing the total number of supported 25-Gbps ports to 80.

  A total of four models are available: two featuring AC power supplies and front-to-back or back-to-front airflow and two featuring DC power supplies and front-to-back or back-to-front airflow.

  [See EX4650 Documentation.]

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **802.1X authentication on trunk ports (EX Series)**—Starting with Junos OS Release 18.3R1, 802.1X authentication can be enabled on trunk ports. Authentication on the trunk port is supported only in single supplicant and single-secure supplicant modes.
- **Multi-domain authentication (EX Series)**—Starting with Junos OS Release 18.3R1, multidomain authentication is supported on EX Series switches. Multidomain authentication is an extension of multiple supplicant mode for 802.1X authentication, and allows one VoIP client and multiple data clients to authenticate to different VLANs while on the same port.

  [See Understanding 802.1X and VoIP on EX Series Switches.]
- **Disable LLDP TLVs (EX2300 and EX3400 switches)**—Starting in Junos OS Release 18.3R1, you can disable specific or all nonmandatory time, length, and value (TLV) messages from being advertised by the Link Layer Discovery Protocol (LLDP) or Link Layer Discovery Protocol—Media Endpoint Discovery (LLDP-MED).

  [See LLDP Overview.]
- **Support for password change policy enhancement (EX Series)**—Starting in Junos OS Release 18.3R1, the Junos password change policy for local user accounts is enhanced to comply with certain additional password policies. As part of the policy improvement, you can configure the following:
  - **minimum-character-changes**—The number of characters by which the new password should be different from the existing password.
  - **minimum-reuse**—The number of older passwords, which should not match the new password.

  [See password.]

Class of Service (CoS)

- **Support for CoS on EX4650 switches (EX4650)**—Starting in Junos OS Release 18.3R1, the EX4650 switch supports CoS functionality. CoS is the assignment of traffic flows to different service levels. You can use CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to ensure quality of service (QoS) to particular applications served by specific traffic flows across the network.

  Compared to CoS functionality on EX4600 switches, EX4650 switches provide significantly more buffer memory (32 MB), but do not support hierarchical scheduling or ETS. The EX4650 also supports eight unicast and two multicast queues.

  [See CoS Support on QFX Series Switches, EX4600 Line of Switches, and QFabric Systems.]

EVPNs

- **EVPN P2MP bud node support (EX9200)**—Starting in Junos OS Release 18.3R1, Junos OS supports configuring a point-to-multipoint (P2MP) label-switched path (LSP) as a provider tunnel on a bud node. The bud node functions both as an egress node and a transit node.

  To enable a bud node to support P2MP LSP, include the **evpn p2mp-bud-support** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level.

  [See Configuring Bud Node Support.]

General Routing

- **Layer 3 unicast features (EX4650)**—Starting with Junos OS Release 18.3R1, the following Layer 3 unicast features are supported:
  - Static routing, ping, and traceroute (IPv4, IPv6)
  - OSPFv2 (IPv4) and OSPFv3 (IPv6)
  - RIPv2
  - BGP (IPv4, IPv6), BGP 4-byte ASN support, and BGP multipath
  - MBGP (IPv4)
  - IS-IS (IPv4, IPv6)
  - BFD (for RIP, OSPF, IS-IS, BGP, PIM)
  - Unicast reverse path forwarding (RPF)
  - Filter based forwarding (FBF)
  - IP directed broadcast traffic forwarding
  - IPv4 over GRE
  - Virtual router redundancy protocol (VRRP)
  - VRRPv3 (IPv6)

- Neighbor Discovery Protocol (IPv6)
- Path MTU discovery
- IPv6 class of service—Behavior aggregate (BA) classifiers, multifield (MF) classifiers and rewrite rules, traffic-class scheduling)
- IPv6 stateless address autoconfiguration
- Equal-cost multipath (ECMP)—32-way
- VXLAN Layer 3 gateway
- MPLS over UDP
- Virtual router (VRF-lite) IS-IS, RIP, OSPF, BGP

Interfaces and Chassis

- **Multichassis link aggregation group (MC-LAG) (EX4650 switches)**—Starting with Junos OS Release 18.3R1, MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

  On one end of an MC-LAG is an MC-LAG client that has one or more physical links in a LAG. This client does not need to detect the MC-LAG. On the other side of the MC-LAG are two MC-LAG switches. Each of these switches has one or more physical links connected to a single client. The switches coordinate with each other to ensure that data traffic is forwarded properly.

  To configure an MC-LAG, include the following statements:

  - **mc-ae** statement at the **[edit interfaces *interface-name* aggregated-ether-options]** hierarchy level
  - **iccp** statement at the **[edit protocols]** hierarchy level
  - **multi-chassis** statement at the **[edit]** hierarchy level

  [See Multichassis Link Aggregation Features, Terms, and Best Practices.]

- **Resilient hashing support for link aggregation groups and equal cost multipath routes (EX4650 switches)**—Starting with Junos OS Release 18.3R1, resilient hashing is supported by link aggregation groups (LAGs) and equal cost multipath (ECMP) sets on EX4650 switches. A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient hashing enhances LAGs by minimizing destination remapping when a new member is added to or deleted from the LAG. Resilient hashing works in conjunction with the default static hashing algorithm. It distributes traffic across all members of a LAG by tracking the flow's LAG member utilization. When a flow is affected by a LAG member change, the Packet Forwarding Engine (PFE) rebalances the flow by reprogramming the flow set table. Destination paths are remapped when a new member is added to or existing members are deleted from a LAG. Resilient hashing applies only to unicast traffic and supports a maximum of 1024 LAGs, with each group having a maximum of 256 members. An ECMP group for a route contains multiple next-hop equal cost addresses for the same destination in the routing table. (Routes of equal cost have the same preference and metric values.)

Junos OS uses a hash algorithm to choose one of the next-hop addresses in the ECMP group to install in the forwarding table. Flows to the destination are rebalanced using resilient hashing. Resilient hashing enhances ECMPs by minimizing destination remapping when a new member is added to or deleted from the ECMP group.

[See Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups.]

- **Channelizing Interfaces on EX4650-48Y Switches**—On the EX4650-48Y switch, there are a total of 56 ports. Of these 56 ports, 8 ports (labeled 48 through 56) are uplink ports that support 100-Gigabit Ethernet interfaces (QSFP28 ports) and 40-Gigabit Ethernet interfaces(QSFP+ ports). The other 48 ports (labeled 0 through 47) are SFP+ ports that support 25-Gigabit Ethernet interfaces or 10-Gigabit Ethernet interfaces. The default speed for the SFP+ ports is 10 Gbps.

  Starting with Junos OS Release 18.3R1, you can channelize the 100-Gigabit Ethernet interfaces to four independent 25-Gigabit Ethernet interfaces. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated 40-Gigabit Ethernet interfaces, or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables on the EX4650-48Y switch.

  > **NOTE:** The uplink ports on the EX4650-48Y switches support auto-channelization.

  If you have disabled auto-channelization, then to channelize the ports, manually configure the port speed using the **set chassis fpc slot-number port port-number channel-speed speed** command, where the speed can be set to 10G or 25G. If a 100-Gigabit Ethernet transceiver is connected, you can only set the speed to 25G. For the SFP+ ports, you can set the speed to 25G or 1G. There is no commit check for this, however.

  > **NOTE:** You cannot configure channelized interfaces to operate as Virtual Chassis ports.

  [See Channelizing Interfaces on Switches.]

Junos Telemetry Interface

- **Routing Engine and Packet Forwarding Engine sensors for the Junos Telemetry Interface (EX4650 and QFX5120-48Y switches)**—Starting with Junos OS Release 18.3R1, Routing Engine and Packet Forwarding Engine statistics are supported through the Junos Telemetry Interface on EX4650 and QFX5120-48Y switches with the same level of support found on QFX5100 switches using Junos OS Release 18.1R1.

  The following Routing Engine statistics are supported through JTI:

  - LACP state export
  - Chassis environmentals export
  - Network discovery chassis and components

- LLDP export and LLDP model
- BGP peer information (RPD)
- RSVP interface export
- RPD task memory utilization export
- LSP event export
- Network Discovery ARP table state
- Network Discovery NDP table state

The following Packet Forwarding Engine statistics are supported through JTI:

- Congestion and latency monitoring
- Logical interface
- Filter
- Physical interface
- LSP
- NPU/LC memory
- Network Discovery NDP table state

Only gRPC streaming is supported.

To provision the sensor to export data through remote procedure call (gRPC), use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **Expanded physical interface queue and traffic statistics sensors for Junos Telemetry Interface (JTI) (PTX, MX, EX, QFX, ACX)**—Starting with Junos OS Release 18.3R1, additional resource paths are added to stream physical (IFD) statistics.

  Prior to Junos OS Release 18.3R1, both traffic and queue statistics for physical interfaces (IFD) are sent out together using the resource path **/interfaces** for gRPC streaming (which is internally used to create **/junos/system/linecard/interface/**) or **/junos/system/linecard/interface/** for UDP (native) sensors.

  Now, traffic and queue statistics can be delivered separately. Doing so can reduce the reap time for non-queue data for platforms supporting Virtual Output Queues (VOQ).

  The following UDP resource paths can be configured:

  - **/junos/system/linecard/interface/** is the existing resource path (no change). Traffic and queue statistics are sent together.
  - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.

- **/junos/system/linecard/interface/queue/** exports queue statistics.

The gRPC resource path **/interfaces** now has the following behavior:

- In releases prior to Junos OS 18.3R1, it delivers all IFD traffic and queue statistics. In Junos OS 18.3R1 and higher, it delivers statistics in two sensors:
  - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
  - **/junos/system/linecard/interface/queue/** exports queue statistics.

To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

[See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

Layer 2 Features

- **Layer 2 unicast features (EX4650 switches)**—Starting with Junos OS Release 18.3R1, the following Layer 2 unicast features are supported:
  - 802.1Q VLAN trunking
  - P-VLAN
  - IRB
  - Layer 3 Vlan-tagged logical interfaces
  - 4096 VLANs
  - MAC address filtering
  - MAC address aging configuration
  - Static MAC address assignment for interfaces
  - Per-VLAN MAC learning (limit)
  - MAC learning disable
  - Persistent MAC (sticky MAC)
  - Q-in-Q tag manipulation
  - MAC address limit per port
  - MAC limiting

- MAC limiting per port, per VLAN
- MAC move limiting
- P-VLAN on Q-in-Q
- 802.1D
- 802.1w (RSTP)
- 802.1s (MST)
- BPDU protection
- Loop protection
- Root protection
- VSTP
- RSTP and VSTP running concurrently
- Link aggregation (static and dynamic) with LACP (fast and slow LACP)
- LLDP
- Multiple VLAN Registration Protocol (802.1ak)

[See Ethernet Switching User Guide.]

- **Layer 2 unicast features (EX4650 switches)**—Starting with Junos OS Release 18.3R1, you can use the Unified Forwarding Table (UFT) feature to allocate forwarding table resources to optimize the memory available for different address types based on the needs of your network. You can choose to allocate a higher percentage of memory for one type of address or another.

  [See Understanding the Unified Forwarding Table.]

MPLS

- **MPLS support (EX4650)**—Starting with Junos OS Release 18.3R1, the following MPLS features are supported:
  - LDP (tunneling over RSVP, targeted LDP, LDP over RSVP)
  - RSVP-TE
  - TE++ container LSPs
  - Automatic bandwidth allocation on LSPs
  - IPv6 tunneling over an MPLS IPv4 network (6PE and 6VPE)
  - Ethernet-over-MPLS (L2 circuit)
  - Layer 3 VPN
  - Carrier-of-carrier VPNs
  - ECMP routing

- Segment routing
- EVPN-VXLAN
- MPLS over IRB interfaces
- VRF support in IRB Interfaces

[See MPLS Feature Support on QFX Series and EX4600 Switches.]

Multicast

- **IGMP snooping with private VLANs (EX4300 switches and EX4300 Virtual Chassis)**—Starting in Junos OS Release 18.3R1, EX4300 switches and EX4300 Virtual Chassis support IGMP snooping with private VLANs (PVLANs). A PVLAN consists of secondary isolated and community VLANs configured within a primary VLAN. Without IGMP snooping support on the secondary VLANs, multicast streams received on a primary VLAN are flooded to the secondary VLANs. This feature extends IGMP snooping on a primary VLAN to its secondary VLANs, which further constrains multicast streams only to interested receivers on PVLANs. When IGMP snooping is enabled on a primary VLAN, it is implicitly enabled on all secondary VLANs, and the secondary VLANs learn the multicast group information on the primary VLAN.

  > **NOTE:** Ports in a secondary VLAN cannot be used as IGMP multicast router interfaces. Secondary VLANs can receive multicast data streams ingressing on promiscuous trunk ports or inter-switch links acting as multicast router interfaces.

  [See IGMP Snooping Overview.]

- **Multicast VLAN registration (MVR) (EX4300 switches and EX4300 Virtual Chassis)**—Starting in Junos OS Release 18.3R1, EX4300 switches and EX4300 Virtual Chassis support multicast VLAN registration (MVR). MVR efficiently distributes IPTV multicast streams across an Ethernet ring-based Layer 2 network, reducing the bandwidth required for this traffic by using a multicast VLAN (MVLAN) over which multicast traffic is forwarded to interested listeners on other VLANs that are configured as MVR receiver VLANs. You can configure MVR at the **[edit protocols igmp-snooping vlan *vlan-name* data-forwarding] source** and **receiver** hierarchy levels, and use the **show igmp snooping data-forwarding** CLI command to view configured MVLAN and MVR receiver VLAN associations.

  [See Understanding Multicast VLAN Registration.]

- **Layer 3 multicast features (EX4650)**—Starting with Junos OS Release 18.3R1, the following Layer 3 multicast features are supported:
  - IGMP version 1 (IGMPv1), version 2 (IGMPv2), and version 3 (IGMPv3)
  - IGMP filtering
  - PIM sparse mode (PIM-SM)
  - PIM dense mode (PIM-DM)

- PIM source-specific multicast (PIM-SSM)
- MSDP

IGMP and PIM are also supported on virtual routers.

[See Multicast Overview.]

- **Layer 2 multicast features (EX4650)**—Starting with Junos OS Release 18.3R1, the following Layer 2 multicast features are supported:
  - IGMP snooping for IGMPv1, IGMPv2, and IGMPv3
  - IGMP proxy
  - IGMP querier

  IGMP snooping is also supported on virtual routers.

  [See Multicast Overview.]

Network Management and Monitoring

- **Customized MIBs for sending custom traps based on syslog events (EX Series)**—Starting in Junos OS Release 18.3R1, there is a process whereby customers can define their own MIBs for trap notifications. The customized MIB maps a particular error message with a custom OID rather than a generic one. Juniper Networks provides two new MIB roots reserved for customer MIBs, one for the custom MIB modules and the other for the trap notifications. For this process, you must convert the MIB to YANG format, and a tool is available for that.

  [See Customized SNMP MIBs for Syslog Traps.]

- **MIB support for media attachment unit (MAU) information (EX2300, EX3400, and EX4300 switches)**—As of Junos OS Release 18.3R1, remote agents can use SNMP to gather information about media attachment units (MAUs) connected to switches. These switches will populate the Entity (RFC 4133) and Entity State (RFC 4268) standard SNMP MIBs and a new MIB table, ifJnxMediaTable, which is part of the Juniper Networks enterprise-specific interface MIB extensions. The objects in the table represent MAU information such as media type, connector type, link mode, and link speed.

  [See SNMP MIB Explorer.]

- **Services support: sFlow, port mirroring, and storm control (EX4650 switches)**—Starting in Junos OS Release 18.3R1, the following services are provided on EX4650 switches:
  - sFlow networking monitoring technology—Collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously.
  - Local and remote port mirroring and remote port mirroring to an IP address—Copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface (local port mirroring), to a VLAN (remote port mirroring), or to the IP address of a device running an analyzer application on a

remote network (remote port mirroring to an IP address [GRE encapsulation]). (When you use remote port mirroring to an IP address, the mirrored packets are GRE-encapsulated.)

- Storm control—Causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

[See Overview of sFlow Technology, Understanding Port Mirroring, and Understanding Storm Control.]

Operation, Administration, and Maintenance (OAM)

- **Connectivity Fault Management (CFM) Support (EX4600)**—IEEE 802.1ag Connectivity Fault Management (CFM) provides fault isolation and detection over large Layer 2 networks which may span several service provider networks. You can configure CFM to monitor, isolate, and verify faults in these interconnected provider bridge networks. Starting in Junos OS Release 18.3R1, Junos OS provides CFM support on EX4600.

  CFM support on EX4600 has the following limitations:

  - CFM support is provided via software using filters. This can impact scaling.
  - Inline Packet Forwarding Engine (PFE) mode is not supported. In Inline PFE mode, you can delegate periodic packet management (PPM) processing to the Packet Forwarding Engine (PFE) which results in faster packet handling and the CCM interval supported is 10 milliseconds.
  - Performance monitoring (ITU-T Y.1731 Ethernet Service OAM) is not supported.
  - CCM interval of less than 1 second is not supported.
  - CFM is not supported on Routed Interfaces and aggregated Ethernet (lag) interfaces.
  - MIP half function, to divide the MIP functionality into two unidirectional segments to improve network coverage, is not supported.
  - Up MEP is not supported.
  - Total number of CFM sessions supported is 30.

[See Understanding Ethernet OAM Connectivity Fault Management for an EX Series Switch.]

Port Security

- **IPv6 Router Advertisement (RA) Guard (EX4600)**—Starting with Junos OS Release 18.3R1 for EX Series switches, IPv6 RA guard is supported on EX4600 switches. RA guard protects against rogue RA messages generated either maliciously or unintentionally by unauthorized or improperly configured routers connecting to the network segment. RA guard works by validating RA messages based on whether they meet certain criteria, which is configured on the switch as a policy. RA guard inspects the RA message and compares the information contained in the message attributes to the policy. Depending on the policy, RA guard either drops or forwards the RA messages that match the conditions.

  [See Understanding IPv6 Router Advertisement Guard.]

Restoration Procedures and Failure Handling

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (EX Series)**—Starting in Junos OS Release 16.1, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode. The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

  [See Saving a Rescue Configuration File.]

Security

- **Support for firewall filters (EX4650)**—Starting with Junos OS Release 18.3R1, you can configure firewall rules to filter incoming network traffic based on a series of user-defined rules. You can specify whether to accept, permit, deny, or forward a packet before it enters an interface. If a packet is accepted, you can also configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). You configure firewall filters at the **[edit firewall]** hierarchy level.

  [See Firewall Filters Overview.]

- **Support for distributed denial-of-service protection (EX4650)**—Starting with Junos OS Release 18.3R1, you can configure denial-of-service (DoS) protection on the switches to continue to function while under attack. A denial-of-service (DoS) attack is any attempt to deny valid users access to network or server resources by using up all the resources of the network element or server. DDoS protection identifies and suppress malicious control packets while enabling legitimate control traffic to be processed. A single point of DDoS protection management enables you to customize profiles for your network control traffic. To protect against DDoS attacks, you can configure policers for host-bound exception traffic. The policers specify rate limits for all control traffic for a given protocol. You can also monitor policers,

obtaining information such as the number of violations encountered and the number of packets received or dropped.

[See Understanding Distributed Denial-of-Service Protection on QFX Series Switches.]

Software Installation and Upgrade

- **Phone-home client (EX4300 switches)**—Starting with Junos OS Release 18.3R1, you can use either the legacy DHCP-options-based ZTP or the phone-home client (PHC) to provision software for the switch. If the switch boots up and there are DHCP options received from the DHCP server for ZTP, ZTP resumes. If DHCP options are not present, PHC is attempted. PHC enables the switch to securely obtain bootstrapping data, such as a configuration or software image, with no user intervention other than having to physically connect the switch to the network. When the switch first boots, PHC connects to a redirect server, which will redirect to a phone home server to get the configuration or software image.

  To initiate either DHCP-options-based ZTP or PCH, the switch must either be in a factory-default state, or you can issue the **request system zeroize** command.

  [See Understanding the Phone-Home Client.]

System Management

- **Secure Boot (EX4650 switches)**—Starting with Junos OS Release 18.3R1, a significant system security enhancement is being introduced: Secure Boot. The secure boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement Secure Boot.

SEE ALSO

Changes in Behavior and Syntax | 49

Known Behavior | 52

Known Issues | 56

Resolved Issues | 61

Documentation Updates | 69

Migration, Upgrade, and Downgrade Instructions | 69

Product Compatibility | 71

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for the EX Series.

Interfaces and Chassis

- **No support for performance monitoring on ae interfaces (EX4300)**—Y.1731 performance monitoring (PM) over aggregated Ethernet interfaces is not supported on EX4300 switches. [See sla-iterator-profile.]

Junos OS XML API and Scripting

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (EX Series)**—Starting in Junos OS Release 18.3R1, Junos OS does not support configuring an MD5 or SHA-1 checksum hash to verify the integrity of local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) scripts or support using an MD5 or SHA-1 checksum hash with the **op url *url* key** option to verify the integrity of remote op scripts.

Layer 2 Features

- **Configuration option for LLDP VLAN name type, length, and value (TLV) (EX3400, EX4300)**—Starting in Junos OS Release 18.3R1, you can configure the **vlan-name-tlv-option (name | vlan-id)** statement at the **[edit protocols lldp]** hierarchy level to select whether to transmit the VLAN name or simply the VLAN ID for the Link Layer Discovery Protocol (LLDP) VLAN name TLV when exchanging LLDP messages. By default, EX Series switches running Enhanced Layer 2 Software (ELS) transmit the VLAN ID for the LLDP VLAN name TLV, and the **show lldp detail** command displays the default string **vlan-*vlan-id*** for an interface's VLAN name in the **Vlan-name** output field. Switches that support the **vlan-name-tlv-option** statement behave the same as the default if you configure the **vlan-id** option with this statement. If you configure the **name** option, the switch transmits the VLAN name instead, and the **show lldp detail** command displays the VLAN name in the **Vlan-name** output field.

Network Management and Monitoring

- **Junos OS does not support management of YANG packages in configuration mode (EX Series)**—Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages using the **run** command in configuration mode is not supported.
- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (EX Series)**—Starting in Junos OS Release 18.3R2, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

Security

- **Firewall warning message (EX2300 switches)**—Starting in 18.3R1, a warning message is displayed whenever a firewall term includes log or syslog with the accept filter action.
- **Syslog or log action on firewall drops packets (EX4600 switches)** —Starting in 18.3R2, if you configure a syslog or log action on an ingress firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (EX Series)**—Starting in Junos OS Release 18.3R1, the jdhcpd process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
  - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
  - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

  The new behavior is as follows:

  1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
  2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

  In earlier releases:

  1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
  2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

  [See Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview.]

Virtual Chassis

- **New configuration option to disable automatic Virtual Chassis port conversion (EX4300 and EX4600 Virtual Chassis)**—Starting in Junos OS Release 18.3R1, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in an EX4300 or EX4600 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches.

When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:

- LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
- The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
- The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

Automatic VCP conversion is not needed when using default-configured VCPs on both sides of the link to interconnect two members. On both ends of the link, you can also manually configure network or uplink ports that are supported as VCPs, whether or not the automatic VCP conversion feature is enabled.

Deleting the **no-auto-conversion** statement from the configuration returns the Virtual Chassis to the default behavior, which reenables automatic VCP conversion.

SEE ALSO

New and Changed Features | 35

Known Behavior | 52

Known Issues | 56

Resolved Issues | 61

Documentation Updates | 69

Migration, Upgrade, and Downgrade Instructions | 69

Product Compatibility | 71

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 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.

Class of Service (CoS)

- On EX4650 switches if the CoS configurations are modified when egress traffic shaped at very low rate (less than 50 Mbps), packets might get stuck in the MMU buffers permanently. It might cause ingress or egress traffic drops. When low rate shapers (less than 50 Mbps) are applied on egress queues, we suggest to deactivate shaping before any CoS modification or ensure traffic is stopped before doing CoS modification. PR1367432

Ethernet Switching

- With software MAC learning enabled, for example, with features such as MAC limiting, MAC move limit, 802.1X authentication, and source MAC filters, MAC learning is slower than with hardware MAC learning. PR1355758

Infrastructure

- Issue is specific to downgrade (17.4T) and core is seen only once during downgrade due to timing issue in sdk toolkit upgradation after which dcpfe recovers by its own and no issues will be seen after that. PR1337008

Layer 2 Features

- For EX4650 the switch might learn its own MAC address on the network interface if it is attached an IRB interface to a VLAN. As a result of the wrong MAC learning, it might result in wrong forwarding in a MC-LAG scenario. PR1365942

Interfaces and Chassis

- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including master routing instance), but only one logical interface was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating incorrect configuration. This issue is fixed and it is now not allowed to configure the same IP address (the length of the mask does not matter) on different logical interfaces. PR1221993

Platform and Infrastructure

- On EX2300 and EX3400 switches, L2PT will not work with tag-protocol-id 0x9100. PR1333475
- Smartd verification is not supported on EX4300-48-MP. Instead, "ssd-stats" can be used from Host-OS to get an overall current health status of SSD. PR1343091

- On EX4300-48MP when primary ROOT Partition is corrupted and switch is power cycled, then switch will get stuck at Linux after boot. Switch needs to be manually rebooted from secondary SSD Partition and recover corrupted primary partition. PR1344938
- Broadcast route is not pingable when NTP is configured in broadcast mode. Ping to Broadcast route is not supported. PR1347480
- **DIRECTORY CORRUPTED I=149350 OWNER=0 MODE=40755** messages continuously printed in console during device boot up after power cycle of the device The error logs are coming from inside Junos VM. As soon as any disk write operation is initiated from inside the VM, it will be written on host disk as well. However, if power cycle happens before disk write completes, this issue is bound to occur. PR1361094
- Logical interfaces statistics are not supported for L2 and aggregated Ethernet interfaces, it is supported only for Layer 3 interfaces (Layer 3 interface should not be member of aggregated Ethernet), please make sure you have only only normal Layer 3 interface. PR1361185
- Bi-directional optics channelization is not supported. PR1361891
- In QFX5000 switches when more than one interface is attached to an output VLAN for remote port mirroring, the traffic will be received by only one of the interfaces. PR1363358
- Few error messages related to function rt_mesh_group_add_check() will be seen during reboot and are harmless. PR1365049
- Auto channelization not supported for 40GBASE-BXSR QSFP+40GE-LX4 QSFP-100G-PSM4 100GBASE-BXSR. PR1366103
- QFX5120/EX4650: with 288k MAC scale, Routing Engine command **show ethernet-switching table summary** output will show the learned scale entries after a delay of around 60 seconds. PR1367538
- Sub-second BFD interval timer is not supported for EX4650 switches. PR1368671
- Since this is Vm based system the recovery would be done from Linux recovery. PR1371014
- Intermittently after JUNOS reboot two of channelized 25G ports using 4x25G breakout cable may not come up. PR1384898
- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. Device can be recovered using power-cycle of the device. PR1385970

Routing Protocols

- Issuing the command "scp -l" in the JUNOS shell, will cause a core fire generation. PR1363973
- Could scale ISISv4, 254 neighbor and 200k routes together. Beyond 200k routes with 254 neighbor, Adjacency flaps and thus traffic drop are noticed. However, with 40 neighbor 351k routes got scaled. PR1368106

- Since the flex counters are shared among IFPs and other tables, in an uni-dimensional testing, ipmc stats counter created will not be equivalent to number of ipmc entries created and stat counter creation will fail with error "No resources for operation" after 60,000 entries. PR1371399
- The mcsnoopd error messages are seen in logs while adding or deleting IGMP PIM configuration. These are debug messages and are not harmful. PR1371662

Virtual Chassis

- A Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur. PR1347902

SEE ALSO

New and Changed Features | 35

Changes in Behavior and Syntax | 49

Known Issues | 56

Resolved Issues | 61

Documentation Updates | 69

Migration, Upgrade, and Downgrade Instructions | 69

Product Compatibility | 71

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 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.

General Routing

- On an EX9200-12QS line card, interfaces with the default speed of 10 Gigabit Ethernet are not brought down even when the remote end of a connection is misconfigured as 40 Gigabit Ethernet. PR1175918
- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, fpc7 KERNEL/PFE APP=NH OUT OF SYNC: **error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none**. No service impact is seen in MPC2 and MPC3 type cards. PR1205593
- Interface range is not supported for channelized interfaces on the EX9253. The user has to configure interfaces individually. PR1350635
- When me0 ports are connected between two EX3400 switches, the link does not come up. The link comes up when me0 is connected to network port. PR1351757
- The working uplink module SFP-T might go down with Junos OS Release 17.2R1 and later releases. PR1360602
- When a VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter will not be installed. PR1362609
- On EX4650 switches, after deleting sFLOW configuration, every five minutes the following error message **sflow_net_socket_init, 423sflow socket connect failed (socket closed)** is displayed on the VTY console. PR1363381
- On EX4650 switches, if lcmd is restarted, a chassisd core file is generated with a traffic drop for a few seconds. PR1363652
- The time lapse between interface down interrupt detection to FRR call back is approximately 148ms on the QFX5120 platform, though the in-place update FRR programming completes in 1ms. The minimum FRR time achieved with this limitation is approximately 150ms and maximum is approximately 275ms. PR1364244
- When an unified ISSU from Junos OS Release 15.1R7.7 to Junos OS Release 16.1R7.6 is performed on an EX9200 Routing Engine, integrated routing and bridging (IRB) IPv4 and IPv6 traffic is dropped. This traffic loss occurs towards the end of the unified ISSU operation when the new backup Routing Engine comes up and synchronizes with the new master Routing Engine. PR1365149
- EX4300 Virtual Chassis systems might fail to register some jnxOperating SNMP OIDs related to the Routing Engines. This behavior is more likely if Virtual Chassis members 0 and 1 (FPC0 and FPC1) are not selected as Routing Engines. PR1368845

- Traffic drop might be observed with a swap out of a Virtual Chassis of QFX5100 to the EX9253 for testing some heavy multicast traffic, even when the IRB interface comes up. PR1369099
- Multicast router advertisement (RA) packets arriving at a VLAN need to be flooded on ports of all FPCs belonging to the same VLAN. Packets when traversing through a HighGig port need to hit the hardware filter to transmit packets in other FPCs. In issue state, the filter is not applicable for the HighGig ports, so multicast RA packets are not traversing through other FPCs. PR1370329
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur; and even with **enhanced convergence** configured there is no guarantee that subsecond convergence will be achieved. PR1371493
- When both flap-on-disconnect and port-bounce are sent, flap-on-disconnect takes precedence, the switch might not trigger link flap. So the device connected to the switch might not initiate DHCP request to allocate an IP address in the new subnet. The CLI command **show dot1x statistics** displays the number of port bounce requests received. PR1372619
- An EX4300 configured with a firewall filter on lo0 and DHCP-security on VLAN simultaneously might drop legitimate DHCP renew requests from clients on the corresponding VLANs. This occurs because of the implementation design and chipset limitation. PR1376454
- After the MACsec session is deleted, the corresponding interfaces might lose their MACsec function when LACP is enabled on them and the statement **exclude lacp** is configured under the **[edit security macsec]** hierarchy. PR1378710
- On EX9200 Series platforms, if there is a packet-length keyword under a firewall filter is applied on the interface egress, the configuration is not committed, because of the commit-check failure. PR1378901
- After unified ISSU from Junos OS Releases 18.1R1, 18.2R1 to 18.3R1, EX9200 32x10-Gigabit SFP interfaces are flapped with error **IFRT: 'IFD add' (opcode 3) failed** on EX9214 MCLAG configuration. PR1384670
- On EX4650, an installation error **rcu_sched self-detected stall on CPU** is seen. PR1384791
- Junos OS might hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. Power-cycle the device to recover it. PR1385970
- For EX4300-48MP switches, active SSD firmware upgrade is supported and a power cycle of the switch is not required after the upgrade. PR1389543
- When **show** command takes a long time to display results, the STP might change states as BPDUs are no longer processed and cause outages. PR1390330
- DCPFE does not come up in some instances of abrupt power-off or power-on of EX4650. Power-cycle of the device or host reboot will recover the device. PR1393554
- Need 1-Gbps speed configuration support on EX9251. PR1400651
- After upgrading to Junos OS Release 18.1R3.3, the following output message is seen continuously: **adt7470_set_pwm**. PR1401709
- On EX4650 platforms, uRPF check in strict mode will not work properly. PR1417546

- EFL license on EX4300-XXMP devices fails to get installed. For example, **{master:0} root@router> request system license add terminal Mar 01 12:03:05 [Type ^D at a new line to end input, enter blank line between each license key] EmergencyJUNOS285602007 aeaqia qmlbjd amrrha 2tcmbr gayaqb ycsbdm mjggim gbastv nzuxaz lsebew 45dfoj xgc3ah fbo6ct 7vv3hl ykp4zq 5g6xch szi7aq 3pek5e vh4myw jdi5wq dxyi3c rkgydi 3crzkr szq terminal:1. EmergencyJUNOS285602007: license not valid for this product add license failed (1 errors) This only affects EFL licenses (AFL is not affected) and -MP EX4300 devices**. PR1421033
- On EX2300, EX3400, EX4300, and EX4600, if igmp-snooping is enabled, multicast traffic might be dropped silently. PR1423556
- I2C read errors are seen when an SFP-T is inserted into a disabled state port configured with **set interface <*> disable** command. PR1423858

Infrastructure

- Junos OS might hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. PR1359339
- When an SNMP poll is performed for the following OIDs, the backup Routing Engine returns the value 6 (6=down) for the FAN and 1 (1=unknown) for the PSUs, even though the FAN and PSUs are up. **Fan: 1.3.6.1.4.1.2636.3.1.13.1.6 PSU: 1.3.6.1.4.1.2636.3.1.13.1.6.2**. For a permanent fix, upgrade the chassis to Junos OS Release 15.1R8 or later. PR1360962
- In a private VLAN (PVLAN) multiple switches scenario, on EX2300, EX3400, EX4300, EX4600, and QFX Series switches (except for QFX10000), after rebooting the device, isolated VLAN traffic received from inter-switch link might be dropped. The configuration **inter-switch-link** statement is used when a PVLAN spans multiple switches. PR1388186
- On EX2300, EX2300-C, and EX2300-MP platforms, if Junos OS is with FreeBSD kernel version 11 with the build date on or after 2019-02-12, the switch might stop forwarding traffic or responding to console. A reboot is required to restore the service. PR1442376

Junos Fusion Enterprise

- On a Junos Fusion Enterprise it might take 6 to 30 seconds for the traffic to converge when on the aggregation device JFE is powered OFF or powered ON. PR1257057
- Power over Ethernet (PoE) over Link Layer Discovery Protocol (LLDP) negotiation is not supported in a Junos Fusion Enterprise (JFE) setup. The issue results in powering up failure when a device makes PoE over LLDP negotiation with the JFE. PR1366106

Layer 2 Features

- On EX2300 and EX3400, if L2PT is configured and the user wants to enable LLDP, then the user needs to configure LLDP individually on the port. The interface all option does not work. There is no functional impact. PR1361114
- On EX2300 and EX3400, while configuring L2PT for tunneling LLDP, the LLDP packets are dropped at the L2PT NNI interface. Issue is seen first time when the configuration is done and recovers with reboot. PR1362173
- **eswd[1200]: ESWD_MAC_SMAC_BRIDGE_MAC_IDENTICAL: Bridge Address Add: XX:XX:db:2b:26:81 SMAC is equal to bridge mac hence don't learn** is seen in syslog every few minutes on ERPS owner. The logs occur during ERPS PDU in ERPS setup. This message can be ignored. PR1372422
- On QFX5000 platform, if storm control is applied on multiple ports, storm control logging might not take effect. PR1401086

Multicast

- IGMP query packets might be duplicated between L2 interfaces with IGMP snooping enabled. PR1391753

Platform and Infrastructure

- IGMPv3 neighborship information is now in synchronization with the kernel entries. PR1317141
- ICMPv6 packets are hitting the dynamic ingress filter with higher priority, thus never reaching an MF or static classifier. PR1388324

Routing Protocols

- On a EX4650 with UFT configuration **num-65-127-prefix-4**, when scaled the greater than 64 prefix IPv6 routes, the command **show pfe route inet6 hw lpm** output will show only a single IPv6 entry but not the scaled entries. PR1369320
- On EX4300 and EX4600 switches, if host destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log or syslog action (for example, filter <> term <> then log/syslog), such packets might not be dropped and reach the Routing Engine unexpectedly. PR1379718
- In a multicast routing scenario using PIM, if configuring a static route with qualified-next-hop for multicast source, the rpd process might crash. This is because qualified-next-hop points to the Gateway Family Data Links (GF_DLI) address which PIM is unable to process, resulting in the crash. PR1408443

Subscriber Access Management

- The authd reuse address quickly before jdhcpd has completely cleaned up the old subscriber, which results in flooding error log. The log such as: **jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815**. PR1402653

SEE ALSO

New and Changed Features | **35**

Changes in Behavior and Syntax | **49**

Known Behavior | **52**

Resolved Issues | **61**

Documentation Updates | **69**

Migration, Upgrade, and Downgrade Instructions | **69**

Product Compatibility | **71**

Resolved Issues

IN THIS SECTION

Resolved Issues: 18.3R2 | **62**

Resolved Issues: 18.3R1 | **65**

This section lists the issues fixed in the Junos OS Release 18.3R2 for the EX Series switches.

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

Resolved Issues: 18.3R2

Authentication and Access Control

- DHCPv6 client is not supported in this release for EX4300-48MP. PR1373691

EVPN

- A few minutes of traffic loss might be observed during recovery from link failure. PR1396597

General Routing

- The Routing Engine Packet Forwarding Engine **out-of-sync** errors might be seen in syslog. PR1232178
- The EX4300-32F MACsec session stays down on 1-Gigabit and 10-Gigabit Ethernet links after certain events, when events are performed with traffic running. PR1299484
- On EX3400 and EX2300 platforms, a redirect message is sent from the switch even when **no-redirect** is set for the specified interface. PR1333153
- The FXPC process might crash after adding or deleting a Q-in-Q VLAN to an interface on EX2300 and EX3400 platforms. PR1334850
- The 40G interfaces might not forward traffic. PR1349675
- FPM board status is missing in SNMP MIB walk result. PR1364246
- OAM Ethernet **connectivity-fault-management** configured on aggregated Ethernet interfaces is not supported but no commit error. PR1367588
- Unable to use Ansible to collect RSI from EX9200. PR1367913
- MAC refresh packet might not be sent out from the new primary link after the RTG failover. PR1372999
- The interface in SFP-T module on EX2300 and EX3400 might be down while its peer connected interface is up. PR1374522
- EX4600VC might not send RIPv2 updates when igmp-snooping is enabled. PR1375332
- The interface AE480 or above might be in STP discarding state on the EX9200 switches. PR1378272
- ARP request packets might be sent out with 802.1Q VLAN tag PR1379138
- All interfaces belonging to certain FPCs might be lost after multiple GRES in Virtual Chassis. PR1379790
- On EX3400 switches, the error messages are seen after applying firewall filter to loopback interface. PR1380544
- The dot1x does not work with Microsoft NPS server. PR1381017
- Constant memory leak might lead to FPC memory exhaustion PR1381527
- Commit error is observed for the first time while loading the **mini-PDT base** configurations. PR1383469
- On the EX4650 switch, occasionally two of the channelized 25-Gigabit Ethernet ports that are using 4x25G breakout cable will not come up after Junos OS reboots. PR1384898

- ARP and ethernet-table entry in pointing to an aggregated Ethernet interface whose state is down. PR1385199
- On EX4300-48MP, the **session-option** stanza under the **[access profile]** hierarchy for EX Series platforms is not applicable. PR1385229
- On EX9200 platforms, the warning message **prefer-status-control-active is used with status-control standby** might be seen whenever you commit an operation. PR1386479
- On EX2300 with Q-in-Q **flexible-vlan-tagging** is unable to obtain DHCP IP for IRB after a reboot/power-cycle. PR1387039
- On EX3400 Virtual Chassis, **Error tvp_status_led_set" and " Error:tvp_optics_diag_eeprom_read** syslog errors are seen. PR1389407
- MAC learning might stop working on some LAG interfaces. PR1389411
- "Input rate pps" is not increased on EX2300-MP uplink ports if the packet is a pure Layer 2 packet like non-etherII or non-EtherSnap. PR1389908
- EX3400VC - When an interface in a Virtual Chassis member switch that is not master, is flapped, IGMP query packets 224.0.0.1 are sent to all the ports of members except the master FPC. PR1393405
- PTP over Ethernet traffic might be dropped when IGMP and PTP TC are configured together. PR1395186
- On EX2300, MAC table is not populated after interface-mode change. PR1396422
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. PR1398398
- After upgrading Junos OS Release 15.1X53 to Junos OS Release 18.2R1.9, the EX3400 cannot learn 30,000 MAC addresses. PR1399575
- The FBF routing-instance instance-type "forwarding" is missed for EX Series (EX3400). PR1400163
- MAC-limit with persistent MAC is not working after reboot. PR1400507
- The authd might crash when you issue the **show network-access requests pending** command during authd restart. PR1401249
- The STP does not work when aggregated interfaces number is ae1000 or above in QFX5110 and QFX5200 and ae480 or above in other QFX Series switches. PR1403338
- The l2cpd might crash if the VSTP **traceoptions** and VSTP VLAN **all** commands are configured. PR1407469
- EX3400 PSU status is still taking "check" status even though PSU module has been removed PR1408675
- The chassisd output power budget is received continuously for 5 seconds without any alarm after upgrading to Junos OS Release 18.1R3. PR1414267
- VXLAN encapsulation next hop (VENH) does not get installed during BGP flap or restart routing. PR1415450

Infrastructure

- IfSpeed and IfHighSpeed erroneously reported as zero on EX2300. PR1326902

Junos Fusion Enterprise

- PoE over LLDP negotiation is not supported on Junos Fusion Enterprise setup. PR1366106
- An error **peer_daemon: bad daemon: scpd** is seen on EX9251 switch running Junos OS Release 18.1R1 and 18.1R2. PR1369646
- Juniper Fusion Enterprise : Cannot login to SD cluster though it is recognized by AD properly. PR1395570
- The l2ald might crash and generate a core file when the **clear ethernet-switching table persistent-learning** command is executed. PR1409403
- Extended ports do not adjust MTU in Junos Fusion Enterprise on VOIP-enabled ports. PR1411179

Layer 2 Features

- RTG MAC refresh packets are sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. PR1389695

Layer 3 Features

- The l2ald might crash when the **clear ethernet-switching table persistent-learning** command is issued. PR1381739

Platform and Infrastructure

- Ping does not go through device after WTR timer expires in ERPS scenario. PR1132770
- On EX4300 switches, in a rare situation the remote interface starts flapping unexpectedly. PR1361483
- Login lockout might never expire because the timestamps of **Lockout start** and **Lockout end** are same. PR1373803
- On EX4300-48MP, unsupported 1 Gigabit optics in the 10 Gigabit uplink module might cause interface traffic to be dropped. PR1374390
- Traffic might be silently discarded with indirect next hop and load balancing. PR1376057
- EX4300 upgrade fails during validation of slax script. PR1376750
- ECMP route installation failure with log messages such as unilist install failure might be observed on EX4300 device. PR1376804
- Packet drops on interface if the statement **gigether-options loopback** is configured. PR1380746
- IRB interface does not turn down when the master Chassis is rebooted or halted. PR1381272
- Traffic loss seen in Layer 2 VPN with GRE tunnel. PR1381740
- On the EX4300 switch, if a loss priority value of high is set for multicast packets by a classifier at the ingress interface, the configuration is overridden by the storm-control filter. PR1382893

- The EX4300 device chooses an incorrect bridge ID as the RSTP Bridge ID. PR1383356
- On EX4300-48MP mixed Virtual Chassis, the Power over Ethernet interface maximum power configuration on a member EX4300 gives an error if the power is configured to be more than 30 W. PR1383717
- Unicast DHCP request get misforwarded to backup RTG link on EX4300-VC. PR1388211
- Layer 3 IP route is destroyed after the Layer 2 next hop is changed. PR1389688
- Continuous log messages get printed in EX4300: **17.4 / MCSNOOPD ICCP Context./var/run/iccpd_control addr /var/run/iccpd_control: Connection refused**. PR1391942
- EX4300 OAM LFM might not work on extended-vlan-bridge interface with native vlan configured PR1399864
- Traffic drop is seen on EX4300 when 10G fiber port is using 1 Gigabit Ethernet SFP optics with autonegotiation enabled. PR1405168

Routing Protocols

- The PPM mode for BFD session in EX4300 is centralized and not distributed by default. PR1361800
- On EX4300-48MP, stale VLAN entries are seen after continuous script run involving split, merge, and reboot. PR1363739
- On EX4650 switches, the output of the **show pfe route summary hw** command shows different scale values for the IPv4 and IPv6 LPM routes rather than the supported scale. PR1366579
- EX4300 might drop incoming IS-IS hello packets when IGMP or MLD snooping is configured. PR1400838
- Sometimes, IGMP snooping might not work. As a workaround, restart multicast snooping process. PR1420921

Subscriber Access Management

- EX4300 line of switches /var showing full **/var/log/dfcd_enc** file grows in size. PR1420921

Resolved Issues: 18.3R1

EVPN

- On EVPN-VXLAN scenarios, a traffic black-hole condition might occur on interfaces that are down, but LACP is up. PR1343515
- Proxy ARP might not work as expected in an EVPN environment. PR1368911

High Availability (HA) and Resiliency

- The Backup Routing Engine might go to db prompt after configuration remove and restore is performed. PR1269383

Infrastructure

- Unable to provide management when the em0 interface of FPC is connected to another FPC Layer 2 interface of the same Virtual Chassis. PR1299385
- Upgrade might fail and the file system might be corrupted if there are blocks in the flash/filesystem. PR1317628
- PFC feature might not work on EX4600. PR1322439
- Archiving dmesg file -**/var/run/dmesg.boot**. PR1327021
- Enabling **mac-move-limit** stops ping on **flexible-vlan-tagging** enabled interface. PR1357742
- Core files are generated when an attempt is made to commit the configuration. PR1376362

Interfaces and Chassis

- On EX4300- Virtual Chassis platforms, the MAC address assigned to an aggregated Ethernet member interface is not the same as that of its parent aggregated Ethernet interface upon master Routing Engine halt. PR1333734
- PoE device does not receive PoE power. PR1345234
- Packets might drop on the ICL of an MC-LAG peer when MC-LAG is up. PR1345316

Layer 2 Features

- The dcpfe/fxpc process might crash when you try to allocate large memory on Packet Forwarding Engines with low memory. PR1362332

Network Management and Monitoring

- On EX4600 platforms, unsupported CLI configurations or **show** commands from the CFM hierarchy or sub-hierarchy are allowed. PR1359052
- CFM: Even after toggling multiple times between baseline and CFM configurations, all 30 CFM sessions are not up. PR1360907

Platform and Infrastructure

- The mismatch of VLAN IDs between an logical interface and VLAN configuration might result in a traffic black-hole condition. PR1259310
- On an EX2300 or EX3400 the bridge ID 02:00:00:00:00:10 is assigned irrespective of base MAC addresses. PR1315633
- Incorrect value of optical power is displayed. PR1326642
- CoS is wrongly applied on Packet Forwarding Engine, leading to egress traffic drop. PR1329141

- When exhausting TCAM table, the filter might be incorrectly programmed. PR1330148
- The FXPC process might crash after adding or deleting a QinQ VLAN to an interface on EX2300/EX3400 platforms. PR1334850
- The configured VOIP VLAN scenario does not work when the P-VLAN is configured as VOIP VLAN. PR1335600
- The device might not learn source MAC addresses, which might be stuck in the Hit Pending state. PR1341518
- MAC source address filter with **accept-source-mac** command does not work if MAC move limit is configured. PR1341520
- On EX4300-MP platforms, the backup Linux cannot be installed first when both SSD partitions are corrupted. PR1342168
- A firewall filter might not be programmed in the Packet Forwarding Engine even though TCAM entries are available. PR1345296
- All the DHCP-Reply or DCHP-Offer packets might be discarded by DHCP snooping if the DHCP snooping is not enabled on that VLAN. PR1345426
- On MPC5, the inline-ka PPP echo requests are not transmitted when the anchor-point is lt-*x*/2/*x* or lt-*x*/3/*x* in pseudowire deployment. PR1345727
- After an EX9200 FPC comes online, the CPU usage on other FPCs might be 100% usage and lead to traffic loss for near 30 seconds. PR1346949
- On EX4300 and EX4600s the VLAN translation feature does not work for the control-plane traffic. PR1348094
- On EX4300 platforms, traffic drop might happen if LLC packets are received with DSAP and SSAP as 0x88 and 0x8e, respectively. PR1348618
- Running RSI through console port might cause system crash and reboot. PR1349332
- On EX2300 or 3400 platforms, L2PT LACP MAC rewrite on a PE device sends duplicate BPDUs to the CE devices. PR1350329
- The transit traffic for ECMP might not work after the EX2300 switch reboots. PR1351418
- On EX4300 platforms (Virtual Chassis and standalone) running Junos OS Release 16.1 and later, a firewall filter with action **then syslog** is unable to send syslog messages to the syslog server. PR1351548
- A high usage chassis alarm in the **/var** partition persists on the EX4300 Virtual Chassis when a file is copied from fpc1 (master) to fpc0 (backup). PR1354007
- The ports using the SFP-T transceiver might continue to be up after system halt. PR1354857
- A commit error is observed if the switch is downgraded from Junos OS Release 18.2 or Release 18.3 to Release 17.3R3. PR1355542

- EX4300-48MP: When DAI and IPSG are configured for many VLANs in one go then DAI Statistics for one interface shows garbage (very large) value. PR1355963
- The FPC stops responding because of a memory leak caused by the VTEP traffic. PR1356279
- On EX2300, EX3400, EX4300-MP platforms in a Virtual Chassis setup, dynamic ARP inspection (DAI) might fail after Virtual Chassis switchover when VSTP is enabled along with **no-mac-table-binding**. PR1359753
- On EX2300, EX3400, EX4300-MP and EX2300-MP platforms used as transit switches, the routed traffic sent out of IRB interfaces uses an old MAC address instead of the configured MAC address for the IRB interface. PR1359816
- On EX2300-MP platforms, a wrong fan count of four is shown, instead of three, in jnxFruName, jnxFilledDescr, and jnxContainersCount. PR1361025
- On EX4300-48MP, the 802.1X protocol subsystem takes a long time to respond to management requests and the following error message is displayed: **the dot1x-protocol subsystem is not responding to management requests.** PR1361398
- A nonexistent fan tray 1 is reported by chassisd on EX2300. PR1361696
- On EX4300-MP switches, MACsec AES-GCM-128-XPN and AES-GCM-256-XPN cipher suites are not supported for mge ports. PR1362035
- Unexpected DCD_PARSE_ERROR_SCHEDULER messages are logged when MS-MPC/MS-MIC is brought offline/online. PR1362734
- Some interfaces cannot be added under the MSTP configuration. PR1363625
- On EX4300 or EX4600 platforms, the l2ald process might crash in an 802.1X scenario. PR1363964
- On EX2300 switches, the **show filter hardware summary** command displays incomplete output. PR1364930
- EX3400 l2cpd crashes when configuring MVRP with Private VLAN and RSTP interface all. PR1365937
- The Packet Forwarding Engine might crash if encounters frequent MAC moves. PR1367141
- Issuing the **request system zeroize** command through noninteractive SSH might not erase the configuration on an EX4300. PR1368452
- Unicast ARP packet loop might be observed in a DAI scenario. PR1370607
- NTP broadcast packets are not forwarded out on L2 ports. PR1371035
- On EX4300 platform with LLDP enabled, LLDP advertisement with incorrect auto-negotiation values might be sent. PR1372966
- BOOTP packets may be dropped if BOOTP-support is not enabled at the global level. PR1373807
- The port access list group does not reallocate TCAM slices properly. PR1375022
- EX4300-48MP: Syslog error ?Error in bcm_port_sample_rate_set(ifl_cmd) : Reason Invalid port. PR1376504

SEE ALSO

New and Changed Features | 35
Changes in Behavior and Syntax | 49
Known Behavior | 52
Known Issues | 56
Documentation Updates | 69
Migration, Upgrade, and Downgrade Instructions | 69
Product Compatibility | 71

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 documentation for the EX Series switches.

SEE ALSO

New and Changed Features | 35
Changes in Behavior and Syntax | 49
Known Behavior | 52
Known Issues | 56
Resolved Issues | 61
Migration, Upgrade, and Downgrade Instructions | 69
Product Compatibility | 71

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 70

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

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 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

SEE ALSO

- New and Changed Features | 35
- Changes in Behavior and Syntax | 49
- Known Behavior | 52
- Known Issues | 56
- Resolved Issues | 61
- Documentation Updates | 69
- Product Compatibility | 71

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

To determine the features supported on EX Series switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at https://apps.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 | 35

Changes in Behavior and Syntax | 49

Known Behavior | 52

Known Issues | 56

Resolved Issues | 61

Documentation Updates | 69

Migration, Upgrade, and Downgrade Instructions | 69

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for Junos Fusion Enterprise. Junos Fusion Enterprise is a Junos Fusion that uses EX9200 switches in the aggregation device role. These release notes describe new and changed features, limitations, and known problems in the hardware and software.

> **NOTE:** For a complete list of all hardware and software requirements for a Junos Fusion Enterprise, including which Juniper Networks devices can function as satellite devices, see Understanding Junos Fusion Enterprise Software and Hardware Requirements .

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

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

> **NOTE:** For more information about the Junos Fusion Enterprise features, see the Junos Fusion Enterprise User Guide.

Release 18.3R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 18.3R2.

Release 18.3R1 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Enterprise in Junos OS Release 18.3R1.

SEE ALSO

Changes in Behavior and Syntax

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 18.3R2 for Junos Fusion Enterprise.

SEE ALSO

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 for Junos Fusion Enterprise.

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

Junos Fusion

- On a Junos Fusion Enterprise, it could take 6 to 30 seconds for the traffic to converge when the aggregation device is powered OFF or powered ON. PR1257057

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 for Junos Fusion Enterprise.

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

Junos Fusion Enterprise

- Power over Ethernet over LLDP (Link Layer Discovery Protocol) negotiation is not supported in Junos Fusion Enterprise. The issue results in failure to power up during PoE over LLDP negotiation. PR1366106

SEE ALSO

New and Changed Features | 72

Changes in Behavior and Syntax | 73

Known Behavior | 74

Resolved Issues | 75

Documentation Updates | 77

Migration, Upgrade, and Downgrade Instructions | 77

Product Compatibility | 82

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in the Junos OS Release 18.3R2 for the Junos Fusion Enterprise.

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

Resolved issues: Release 18.3R2

- The scpd process is not running in EX9251, causing an error message in the CLI. PR1369646
- Cannot login to the satellite device cluster on a Junos Fusion Enterprise even though it is recognized by the aggregation device. PR1395570
- The l2ald process might generate a core file if the **clear ethernet-switching table persistent-learning** command is issued. PR1409403
- Extended ports in Junos Fusion Enterprise do not adjust MTU when VoIP is enabled. PR1411179

Resolved issues: Release 18.3R1

- A satellite device does not recover Power over Ethernet after the device is offline for more than 10 minutes and rejoins the aggregation device. PR1356478
- The Fusion satellite device reboots post automatic POE firmware upgrade. PR1359065
- The ppm-lite process might generate a core file on the Fusion satellite devices. It is unexpectedly treating IEEE PORT VLAN ID TLV on LLDP packets as a DCBXv1.01 TLV. PR1364265
- The scpd process is not running in EX9251, causing an error message in the CLI. PR1369646

SEE ALSO

New and Changed Features | 72

Resolved Issues | 75

Known Behavior | 74

Known Issues | 75

Documentation Updates | 77

Migration, Upgrade, and Downgrade Instructions | 77

Product Compatibility | 82

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 for Junos Fusion Enterprise documentation.

SEE ALSO

- New and Changed Features | 72
- Changes in Behavior and Syntax | 73
- Known Behavior | 74
- Known Issues | 75
- Resolved Issues | 75
- Migration, Upgrade, and Downgrade Instructions | 77
- Product Compatibility | 82

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade or downgrade Junos OS and satellite software for a Junos Fusion Enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos Fusion Enterprise topology.

Basic Procedure for Upgrading Junos OS on an Aggregation Device

When upgrading or downgrading Junos OS for an aggregation device, always use the **junos-install** 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 **junos-install** package and details of the installation process, see the Installation and Upgrade Guide.

> **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. See the Junos OS Software Installation and Upgrade Guide.

To download and install Junos OS:

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

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

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

3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list on the right of the page.

5. Select the **Software** tab.

6. Select the software package for the release.

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 **junos-install** package on the aggregation device.

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

Customers in the United States and Canada, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-18.3R2.n.tgz
```

All other customers, use the following commands, where *n* is the spin number.

```
user@host> request system software add validate reboot
source/junos-install-ex92xx-x86-64-18.3R2.n-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to minimize disrupting network operations 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*.

Preparing the Switch for Satellite Device Conversion

There are multiple methods to upgrade or downgrade satellite software in your Junos Fusion Enterprise. See Configuring or Expanding a Junos Fusion Enterprise.

For satellite device hardware and software requirements, see Understanding Junos Fusion Enterprise Software and Hardware Requirements.

Use the following command to install Junos OS on a switch before converting it into a satellite device:

```
user@host> request system software add validate reboot source/package-name
```

> **NOTE:** The following conditions must be met before a Junos switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:
>
> - The Junos switch can only be converted to SNOS 3.1 and higher.
> - The Junos switch must be either set to factory default configuration to factory default configuration using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

When the interim installation has completed and the switch is running a version of Junos OS that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device using the console port.
2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

> **NOTE:** The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port *port-number***

   For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

   `user@satellite-device>`**request virtual-chassis vc-port delete pic-slot 1 port 0**
   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 1**
   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 2**
   `user@satellite-device>` **request virtual-chassis vc-port delete pic-slot 1 port 3**

   This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos Fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, or preconfiguration. See Configuring or Expanding a Junos Fusion Enterprise for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Switch

In the event that you need to convert a satellite device to a standalone device, you will need to install a new Junos OS software package on the satellite device and remove it from the Junos Fusion topology. For more information, see Converting a Satellite Device to a Standalone Device.

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 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

Downgrading Junos OS

Junos Fusion Enterprise is first supported in Junos OS Release 16.1, although you can downgrade a standalone EX9200 switch to earlier Junos OS releases.

> **NOTE:** You cannot downgrade more than three releases.
>
> For more information, see the Installation and Upgrade Guide.

To downgrade a Junos Fusion Enterprise from Junos OS, follow the procedure for upgrading, but replace the **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

- New and Changed Features | 72
- Changes in Behavior and Syntax | 73
- Known Behavior | 74
- Known Issues | 75
- Resolved Issues | 75
- Documentation Updates | 77
- Product Compatibility | 82

Product Compatibility

IN THIS SECTION

- Hardware and Software Compatibility | 83
- Hardware Compatibility Tool | 83

Hardware and Software Compatibility

For a complete list of all hardware and software requirements for a Junos Fusion Enterprise, including which Juniper Networks devices function as satellite devices, see Understanding Junos Fusion Enterprise Software and Hardware Requirements in the Junos Fusion Enterprise User Guide.

To determine the features supported in a Junos Fusion, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: https://apps.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 Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the Junos Fusion Provider Edge. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in the Junos OS main release and the maintenance releases for Junos Fusion Provider Edge.

Release 18.3R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.3R2.

Release 18.3R1 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.3R1.

SEE ALSO

Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 18.3R2.

SEE ALSO

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 for Junos Fusion Provider Edge.

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

SEE ALSO

Known Issues

There are no known issues in the Junos OS Release 18.3R2 for Junos Fusion Provider Edge.

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

SEE ALSO

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in the Junos OS main release and the maintenance releases for Junos Fusion Provider Edge.

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

Resolved Issues: 18.3R2

Junos Fusion Provider Edge

- Laser receive power of extended ports is higher than the output power of the peer link. PR1358007
- Broadcast,Unknown Unicast and Multicast(BUM) traffic might get dropped on peer Fusion Aggregation Device when link between Satellite Device and local Aggregate Device goes down. PR1384440

Junos Fusion Satellite Software

- The shutdown of the cascade port might lead to the invalidation of the MPC linecard. PR1360876
- Extended Port (EP) LAG might go down on the Satellite Devices (SDs) if the related Cascade Port (CP) links to an Aggregation Device (AD) goes down. PR1397992

Resolved Issues: 18.3R1

Junos Fusion

- In Junos Fusion, the aggregation device LAG interface might flap during satellite device upgrade or downgrade. PR1321575
- ppmd crash after changing the mode of EX4300 from standalone to SD. PR1375647
- The spmd core might be seen after executing **request support information** on Aggregation Device. PR1375732

SEE ALSO

New and Changed Features | **84**

Changes in Behavior and Syntax | **85**

Known Behavior | **85**

Known Issues | **86**

Documentation Updates | **88**

Migration, Upgrade, and Downgrade Instructions | **88**

Product Compatibility | **97**

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 documentation for Junos Fusion Provider Edge.

SEE ALSO

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for Junos Fusion Provider Edge. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

Basic Procedure for Upgrading an Aggregation Device

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.

> **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. See the Junos OS Administration Library.

The download and installation process for Junos OS Release 18.3R2 is different from that for earlier Junos OS releases.

1. Using a Web browser, navigate to the Download Software URL on the Juniper Networks webpage:

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

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

3. Select **By Technology > Junos Platform > Junos Fusion** to find the software that you want to download.

4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the page.

5. Select the **Software** tab.

6. Select the software package for the release.

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 aggregation device.

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

Customers in the United States and Canada, use the following commands.

- For 64-bit software:

> **NOTE:** We recommend that you see 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

```
user@host> request system software add validate reboot
source/jinstall64-18.3R2.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-18.3R2.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

> **NOTE:** We recommend that you see 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

```
user@host> request system software add validate reboot
source/jinstall64-18.3R2.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-18.3R2.SPIN-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 the Canada and U.S. version)

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

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

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 18.3R2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

If the aggregation device has two Routing Engines, perform a Junos OS installation on each Routing Engine separately as follows to minimize disrupting network operations:

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*.

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos Fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see Understanding Junos Fusion Software and Hardware Requirements

> **NOTE:** The following conditions must be met before a standalone switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:
>
> - The switch can be converted to only SNOS 3.1 and later.
> - Either the switch must be set to factory-default configuration by using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

Customers with EX4300 switches, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-ex-4300-14.1X53-D43.3-domestic-signed.tgz
```

Customers with QFX5100 switches, use the following command:

```
user@host> request system software add reboot
source/jinstall-qfx-5-14.1X53-D43.3-domestic-signed.tgz
```

When the interim installation has completed and the switch is running a version of Junos OS that is compatible with satellite device conversion, perform the following steps:

1. Log in to the device by using the console port.

2. Clear the device:

```
[edit]
user@satellite-device# request system zeroize
```

> **NOTE:** The device reboots to complete the procedure for resetting the device.

If you are not logged in to the device by using the console port connection, your connection to the device is lost after you enter the **request system zeroize** command.

If you lose your connection to the device, log in using the console port.

3. (EX4300 switches only) After the reboot is complete, convert the built-in 40-Gbps QSFP+ interfaces from Virtual Chassis ports (VCPs) into network ports:

   ```
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port port-number
   ```

   For example, to convert all four built-in 40-Gbps QSFP+ interfaces on an EX4300-24P switch into network ports:

   ```
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 0
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 1
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 2
   user@satellite-device> request virtual-chassis vc-port delete pic-slot 1 port 3
   ```

   This step is required for the 40-Gbps QSFP+ interfaces that will be used as uplink interfaces in a Junos Fusion topology. Built-in 40-Gbps QSFP+ interfaces on EX4300 switches are configured into VCPs by default, and the default settings are restored after the device is reset.

After this initial preparation, you can use one of three methods to convert your switches into satellite devices—autoconversion, manual conversion, and preconfiguration. See Configuring Junos Fusion Provider Edge for detailed configuration steps for each method.

Converting a Satellite Device to a Standalone Device

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove the satellite device from the Junos Fusion topology.

> **NOTE:** If the satellite device is a QFX5100 switch, you need to install a PXE version of Junos OS. The PXE version of Junos OS is software that includes *pxe* in the Junos OS package name when it is downloaded from the Software Center—for example, the PXE image for Junos OS Release 14.1X53-D43 is named install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz . If the satellite device is an EX4300 switch, you install a standard jinstall-ex-4300 version of Junos OS.

The following steps explain how to download software, remove the satellite device from Junos Fusion, and install the Junos OS software image on the satellite device so that the device can operate as a standalone device.

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

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

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

3. Select **By Technology > Junos Platform > Junos Fusion** from the drop-down list and select the switch platform series and model for your satellite device.

4. Select the Junos OS Release 14.1X53-D30 software image for your platform.

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

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

8. Remove the satellite device from the automatic satellite conversion configuration.

   If automatic satellite conversion is enabled for the satellite device's member number, remove the member number from the automatic satellite conversion configuration. The satellite device's member number is the same as the FPC slot ID.

   ```
   [edit]
   user@aggregation-device# delete chassis satellite-management auto-satellite-conversion satellite member-number
   ```

   For example, to remove member number 101 from Junos Fusion:

   ```
   [edit]
   user@aggregation-device# delete chassis satellite-management auto-satellite-conversion satellite 101
   ```

   You can check the automatic satellite conversion configuration by entering the **show** command at the **[edit chassis satellite-management auto-satellite-conversion]** hierarchy level.

9. Commit the configuration.

   To commit the configuration to both Routing Engines:

   ```
   [edit]
   user@aggregation-device# commit synchronize
   ```

Otherwise, commit the configuration to a single Routing Engine:

```
[edit]
user@aggregation-device# commit
```

10. Install the Junos OS software on the satellite device to convert the device to a standalone device.

    ```
    [edit]
    user@aggregation-device> request chassis satellite install URL-to-software-package fpc-slot member-number
    ```

    For example, to install a PXE software package stored in the **/var/tmp** directory on the aggregation device onto a QFX5100 switch acting as the satellite device using FPC slot 101:

    ```
    [edit]
    user@aggregation-device> request chassis satellite install
    /var/tmp/install-media-pxe-qfx-5-14.1X53-D43.3-signed.tgz fpc-slot 101
    ```

    For example, to install a software package stored in the **var/tmp** directory on the aggregation device onto an EX4300 switch acting as the satellite device using FPC slot 101:

    ```
    [edit]
    user@aggregation-device> request chassis satellite install
    /var/tmp/jinstall-ex-4300-14.1X53-D30.3-domestic-signed.tgz fpc-slot 101
    ```

    The satellite device stops participating in the Junos Fusion topology after the software installation starts. The software upgrade starts after this command is entered.

11. Wait for the reboot that accompanies the software installation to complete.

12. When you are prompted to log back into your device, uncable the device from the Junos Fusion topology. See *Removing a Transceiver from a QFX Series Device* or *Remove a Transceiver*, as needed. Your device has been removed from Junos Fusion.

    NOTE: The device uses a factory-default configuration after the Junos OS installation is complete.

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 18.3R2, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

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 17.1, 17.2, and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

Downgrading from Junos OS Release 18.3

To downgrade from Release 18.3 to another supported release, follow the procedure for upgrading, but replace the 18.3 **jinstall** package with one that corresponds to the appropriate release.

> **NOTE:** You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

SEE ALSO

- New and Changed Features | **84**
- Changes in Behavior and Syntax | **85**
- Known Behavior | **85**
- Known Issues | **86**
- Resolved Issues | **86**
- Documentation Updates | **88**
- Product Compatibility | **97**

Product Compatibility

IN THIS SECTION

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. See the 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 | 84
- Changes in Behavior and Syntax | 85
- Known Behavior | 85
- Known Issues | 86
- Resolved Issues | 86
- Documentation Updates | 88
- Migration, Upgrade, and Downgrade Instructions | 88

Junos OS Release Notes for MX Series 5G Universal Routing Platform

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the MX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

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

Release 18.3R2 New and Changed Features

MPLS

- **Control transport address used for targeted-LDP session (MX Series)**—Currently, only the router-ID or interface address is used as the LDP transport address. Starting in Junos OS Release 18.3R2, you can configure any other IP address as the transport address of targeted LDP sessions, session-groups, and interfaces. This new configuration is applicable only for configured LDP neighbors that have Layer 2 circuit, MPLS, and VPLS adjacencies.

  This feature is beneficial when you have multiple loopback interface addresses, and different IGPs associated with LDP interfaces, and you can control the session established between targeted LDP neighbors with the configured transport address.

  [See Control Transport Address Used for Targeted-LDP Session.]

Network Management and Monitoring

- **New major alarms on MX Series routers with MPC1 and MPC2**—Starting in Junos OS Release 18.3R2, on MX Series routers with MPC1 and MPC2 line cards, a major chassis alarm is raised when the following transient hardware errors occur:

  - CPQ Sram parity error
  - CPQ RLDRAM double bit ECC error

  In the **Description** column of **show chassis alarm** outputs, these errors are described as 'FPC <slot number> Major Errors'. See an example below:

  `user@host>` **show chassis alarms**

```
5 alarms currently active
Alarm time                            Class   Description
2018-10-05 18:48:06 PDT               Major   FPC 9 Major Errors
```

  By default, these errors result in the Packet Forwarding Engine interfaces on the FPC being disabled. You can use the **show chassis fpc errors** command to view the default or user-configured action that resulted from the error.

  You can check the syslog messages to know more about the errors. See the following examples:

```
Oct  5 15:58:02  codeine fpc1 MQCHIP(0) CPQ RLDRAM double bit ECC error, bank 0
 addr 0x0
Oct  5 15:58:02  codeine fpc1 MQCHIP(0) CPQ Sram parity error, errlog 0x0
```

To resolve the error, restart the line card. If the error is still not resolved, open a support case using the Case Manager link at https://www.juniper.net/support/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).

Routing Protocols

- **Support for creating IS-IS topology independent LFA for prefix-SIDs learned from LDP mapping server**—Starting in Junos OS Release 18.3R2, you can configure a point of local repair to create a topology independent loop-free alternate backup path for prefix-SIDs derived from LDP mapping server advertisements in an IS-IS network. In a network configured with segment routing, IS-IS uses the LDP mapping server advertisements to derive prefix-SIDs. LDP Mapping server advertisements for IPv6 are currently not supported.

  To attach flags to LDP mapping server advertisements, include the **attached** statement at the **[edit routing-options source-packet-routing mapping-server-entry *mapping-server-name*]** hierarchy level.

Release 18.3R1 New and Changed Features

Hardware

- **Support for JNP-SFP-10G-BX10D and JNP-SFP-10G-BX10U transceivers (MX80, MX104, MX240, MX480, and MX960 with MIC-MACSEC-20GE)**—Starting in Junos OS Release 18.3R1, the MX80, MX104, MX240, MX480, and MX960 installed with the MIC-MACSEC-20GE support the JNP-SFP-10G-BX10D and the JNP-SFP-10G-BX10U transceivers. The JNP-SFP-10G-BX10D and JNP-SFP-10G-BX10U transceivers are for single SMF bidirectional applications. A JNP-SFP-10G-BX10D transceiver should always be connected to a JNP-SFP-10G-BX10U transceiver with a single SMF. The operating link distance is up to 10 km. With a single LC receptacle, the JNP-SFP-10G-BX10D transmits a 1330 nm wavelength signal and receives a 1270 nm signal, whereas the JNP-SFP-10G-BX10U transmits a 1270 nm wavelength signal and receives a 1330 nm signal.

  [See the Hardware Compatibility Tool.]

- **Support for 10-Gbps ports to operate at 1-Gbps speed (MX204 and MX10003)**—Starting in Junos OS Release 18.3R1, you can use the Mellanox 10-Gbps pluggable adapter (QSFP+ to SFP+ adapter or QSA; model number: MAM1Q00A-QSA) to convert 4 lane-based ports to a single lane-based SFP+ port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ module. Use the QSA adapter to convert a 40-Gbps port to a 10-Gbps or a 1-Gbps port.

  NOTE:
  - The interface name prefix must be **xe**.
  - On the MX10003 router, the MACsec MIC does not provide 1-Gbps speed.
  - On MX204 and MX10003 routers, rate selectability at PIC level and port level does not support 1-Gbps speed.

To configure an interface to operate in the 1-Gbps mode, use the **set interfaces xe-0/1/0 gigether-options speed 1g** command at the [edit] hierarchy level.

[See MX10003 MPC Rate-Selectability Overview and MX204 Router Rate-Selectability Overview.]

- **New MIC-MACSEC-20GE (MX80, MX104, MX240, MX480, and MX960)**—Starting with Junos OS Release 18.3R1, MIC-MACSEC-20G is supported on MX80, MX104, MX240, MX480, and MX960 routers. The MIC has 20 SFP ports supporting 20 SFP pluggable optics modules operating in 1-Gbps mode or two SFP+ ports supporting 2 SFP+ pluggable optics modules operating in 10-Gbps mode. The MIC enables resiliency support and MACsec capability on 1-Gbps and 10-Gbps ports on MX80, MX104 and on the MPC1, MPC2, MPC3, MPC2E, MPC3E, MPC2E-NG, and MPC3E-NG line cards of MX240, MX480 and MX960 routers. The resiliency support includes software support to handle hardware failures of various components of the MIC.

  [See Gigabit Ethernet MIC with 256b-AES MACsec].

  NOTE:
  - FPCs in the MX240, MX480, MX960 routers and the FEB in the MX80 and MX104 routers undergo an automatic bounce or a reboot when the speed toggles between 1-Gbps and 10-Gbps.
  - Rate selectability is supported at the PIC level and not at the port level.
  - When the MIC is operating in the 10-Gbps mode, all the other 1-Gbps ports are disabled.

- **QFX-SFP-1GE-T**—Starting with Junos OS Release 18.3R1, the QFX-SFP-1GE-T transceiver is supported on the SFP+ ports on MX204 routers. When using the QFX-SFP-1GE-T transceiver, keep the following limitations in mind:
  - Speed values less than 1 Gbps are not supported.
  - Configuring the speed as **speed 1G** is required and the **no-auto-negotiation** option is not supported. By default, auto-negotiation is enabled.
  - Link aggregation group (LAG) and timing (1588/SyncE) features are not supported.

  See the [Hardware Compatibility Tool]

Authentication, Authorization, and Accounting

- **Support for password change policy enhancement (MX Series)**—Starting in Junos OS Release 18.3R1, the Junos password change policy for local user accounts is enhanced to comply with certain additional password policies. As part of the policy improvement, you can configure the following:
  - **minimum-character-changes**—The number of characters by which the new password should be different from the existing password.
  - **minimum-reuse**—The number of older passwords, which should not match the new password.

See password

- **MD5 is not supported as an authentication encryption mechanism (MX Series)**—Starting with Junos OS Release 18.3R1, the **md5** option at the **[edit system login password]** hierarchy level for user authentication is not supported. However, the **sha1**, **sha256**, and **sha512** options are supported.

  > **NOTE:** Before Junos OS upgrade, if the device configuration contains usernames whose plain text passwords are MD5 encrypted, after upgrade of Junos OS, users can still log in with the same username and plain text password.

EVPN

- **NSR and unified ISSU support for point-to-multipoint LSP for EVPN provider tunnel (MX Series and vMX)**—Starting in Junos OS Release 18.3R1, Junos OS provides nonstop routing (NSR) and unified ISSU support for point-to-multipoint (P2MP) inclusive provider tunnels. This ensures that broadcast, unknown unicast, and multicast (BUM) packets continue after a Routing Engine switchover occurs when NSR is enabled.

  > **NOTE:** Unified ISSU is not supported on the vMX routers.

  [See Understanding P2MPs LSP for the EVPN Inclusive Provider Tunnel.]

- **Support for mLDP P2MP tunnels with EVPN for BUM traffic (MX Series and vMX)**—Starting in Junos OS Release 18.3R1, Junos OS provides the ability to configure and signal a P2MP LSP for the EVPN Inclusive Provider Tunnel for BUM traffic. The P2MP LSP manages efficient core bandwidth utilization as it uses multicast replication only at the required nodes instead of ingress replication at the ingress PE device. The new configuration is added to enable P2MP tunnels for EVPN at the **[edit routing-instances *routing-instance-name* provider-tunnel]** hierarchy level.

  [See Understanding P2MPs LSP for the EVPN Inclusive Provider Tunnel.]

- **EVPN P2MP bud router support (MX Series and vMX)**—Starting in Junos OS Release 18.3R1, Junos OS supports configuring a point-to-multipoint (P2MP) label-switched path (LSP) as a provider tunnel on a bud router. The bud router functions both as an egress router and a transit router.

  To enable a bud router to support P2MP LSP, include the **evpn p2mp-bud-support** statement at the **[edit routing-instances routing-instance-name protocols evpn]** hierarchy level.

  [See Configuring Bud Node Support.]

- **Support for pseudowire termination at an EVPN (MX Series)**—Starting in Junos OS Release 18.3R1, Junos OS supports port-based and VLAN-based pseudowire that terminates at an EVPN on a single-homed PE device. The pseudowire tunnel originates on an aggregation node, carrying VLAN traffic from different

access nodes and supports packets with no VLAN tags, as well as packets with single VLAN tags and dual VLAN tags.

[See Overview of Psuedowire in EVPN.]

- **Connectivity fault management support for MIP in an EVPN with ETREE and ELAN Services and up MEP in EVPN with ETREE services (MX Series)**—Starting with Junos OS Release 18.3R1, Junos OS supports maintenance association intermediate point (MIP) in an EVPN with ELAN and EVPN ETREE services and connectivity fault management (CFM) up maintenance association end points (MEPs) on attachment circuits (ACs) in an EVPN with ETREE services. This feature also supports Ethernet loopback and Ethernet linktrace for a MEP and delay measurement and synthetic loss measurement for performance monitoring between two MEPs. Monitoring is only supported between a leaf and root node or between two root nodes in an EVPN with ETREE services.

  [See Connectivity Fault Management Support for EVPN.]

- **Support for pseudowire termination at an EVPN using RLT (MX Series)**—Starting in Junos OS Release 18.3R1, you can configure a pseudowire tunnel termination at an EVPN using a redundant logical tunnel (RLT). The RLT provides redundancy to the pseudowire tunnel with two logical interfaces, where only one interface is active at any given time. The active and standby logical interface provides redundancy in case of FPC failure.

  [See Overview of Psuedowire in EVPN.]

Forwarding and Sampling

- **Improved hash computation for IPv6 and multiservice flows (MX Series routers with Trio MPCs)**—Starting in Junos OS Release 18.3R1, to improve load balancing in certain cases, the default behavior for calculating the enhanced-hash-key at the **[forwarding-options enhanced-hash-key family inet6]** hierarchy level now includes the **flow-label** field. This hash is used when choosing an ECMP path where there is an aggregate interface.

  Likewise, for **forwarding-options enhanced-hash-key family multiservice**, the default calculation now includes the **payload** field. To revert to the previous method, specify **no-payload**, or **no-flow-label**, in the intended hierarchy.

  [See family multiservice.]

High Availability and Resiliency

- **ARP stability enhancement during ISSU (MX Series)**—Starting in Junos OS Release 18.3R1, Address Resolution Protocol (ARP) entries on MX Series Routers will not expire while ISSU is underway. This prevents issues with ARP renew packets being dropped during ISSU and causing traffic loss.

  [See Getting Started with Unified In-Service Software Upgrade.]

- **Support for MX-VC BBE configurations and ISSU (MX10003)**—Starting with Junos OS Release 18.3R1, MX10003 routers support MX Virtual Chassis (MX-VC) Broadband Edge (BBE) features, including ISSU.

  MX10003 routers can only interoperate with other MX10003 routers in a MX-VC configuration.

[See Unified ISSU in a Virtual Chassis]

Interfaces and Chassis

- **Support for flexible tunnel interfaces (MX Series)**—Starting in Junos OS Release 18.3R1, a new type of interface, called flexible tunnel interface (FTI), is supported on MX Series routers. FTIs support Layer 3 point-to-point tunnels. These tunnels use Virtual Extensible LAN (VXLAN) encapsulation with a Layer 2 pseudo-header. To configure FTIs on your device and to enable multiple encapsulations on the FTIs, use the **vxlan-*gpe*** parameter under the mandatory **tunnel-endpoint *vxlan*** encapsulation at the **[edit interfaces *interface-name* unit *logical-unit-number* tunnel encapsulation**] hierarchy level.

- **Support for PTP over Ethernet and hybrid mode over link aggregation group (MX240, MX480, MX960, MX2010, MX2020)**—Starting in Junos OS Release 18.3R1, the MPC7E, MPCE8E, and MPC9E line cards support Precision Time Protocol (PTP) over Ethernet and hybrid mode over a link aggregation group (LAG).

  Link aggregation is a mechanism of combining multiple physical links into a single virtual link to achieve linear increase in bandwidth and to provide redundancy in case a link fails. The virtual link is referred to as an aggregated Ethernet interface or a LAG.

  [See Precision Time Protocol Overview.]

- **Support for MIC-MACSEC-20GE (MX80, MX104, MX240, MX480, and MX960)**—Starting in Junos OS Release 18.3R1, MIC-MACSEC-20GE, a Media Access Control Security (MACsec)-enabled MIC, is supported on MX80 and MX104 routers and on the MPC1, MPC2, MPC3, MPC2E, MPC3E, MPC2E-NG, and MPC3E-NG MPCs on the MX240, MX480, and MX960 routers. On this MIC, you can configure either twenty 1-Gigabit Ethernet ports or two 10-Gigabit Ethernet ports that support SFP transceivers.

  The 1-Gigabit Ethernet and 10-Gigabit Ethernet port types on MIC-MACSEC-20GE support both 256-bit AES encryption as well as 128-bit AES encryption.

  [See Multi-Rate Ethernet MIC.]

- **Support for SSD upgrade on backup Routing Engines(MX Series)**—Starting in Junos OS Release 18.3R1, you can upgrade the SSD firmware on the backup Routing Engines, RE-S-X6-64G and RE-MX2K-X8-64G Routing Engines, without switching mastership. In releases before Junos OS Release 18.3R1, SSD upgrade is only supported on the master Routing Engine and, to upgrade firmware on the backup Routing Engine, you must switch mastership by using the **request chassis routing-engine master switch** command and then log in to the backup Routing Engine.

  [See Upgrading the SSD Firmware on Routing Engines with VM Host Support.]

- **BGP Monitoring Protocol can run in a non-default management instance (MX Series)**—Starting in Junos OS Release 18.3R1, the BGP Monitoring Protocol (BMP) can send monitoring packets to BMP monitoring stations that are reachable through a VRF. This feature can be used with the **management-instance** configuration statement to create the routing instance mgmt-junos for BMP to move through. Previously, BMP could only send monitoring packets to a BMP monitoring station that could be looked up using the default (inet.0 or inet6.0) routing table.

  [See Configuring BGP Monitoring Protocol to Run Over a Different Routing Instance.]

IPv6

- **ARP and neighbor discovery command enhancements (MX Series)**—Starting with Junos OS Release 18.3R1, enhancements are made to ARP and neighbor discovery command summaries. ARP and Neighbor Discovery protocol (NDP) are used to resolve next hop entries and to maintain next-hop entries in ARP and ND cache.

  The following enhancements are made to the **show arp**, **show ipv6 neighbors**, and **clear ipv6 neighbors** commands:

```
user@host> show arp ?
Possible completions:
  <[Enter]>            Execute this command
  expiration-time      Show seconds remaining before expiration
  hostname             Name of host
  interface            Name of the interface
  logical-system       Name of logical system
  no-resolve           Don't attempt to print addresses symbolically
  vpn                  Name of VPN routing table
 +ref-count           Next hop Reference count
 +state                Arp entry state
```

```
user@host> show ipv6 neighbors ?
Possible completions:
  <[Enter]>            Execute this command
 +host                 Neighbor host IPV6 address
 +interface            Name of the interface
 +logical-system       Name of logical system
 +reference-count      Next hop reference count
 +vpn                  Name of VPN routing table
 +flags                Next hop flags
  |                    Pipe through a command
```

```
user@host> clear ipv6 neighbors ?
Possible completions:
  <[Enter]>            Execute this command
  all                  Clear all neighbors
  host                 Neighbor host IPV6 address
 +interface            Name of the interface
 +logical-system       Name of logical system
 +vpn                  Name of VPN routing table
  |                    Pipe through a command
```

> **NOTE:** These command summaries have the existing parameters along with the additional parameters.

[See show arp, show ipv6 neighbors, and clear ipv6 neighbors.]

Junos Telemetry Interface

- **Support for the Junos Telemetry Interface (ACX6360, MX Series, and PTX Series)**—Starting with Junos OS Release 18.3R1, Junos Telemetry Interface support is available for the ACX6360 Universal Metro Router and MX Series and PTX Series routers with a CFP2-DCO optics module that provides a high-density, long-haul optical transport network (OTN) transport solution with MACsec capability.

  You can provision sensors to export telemetry data to an outside collector.

  The following native (UDP) and gRPC sensors can be provisioned for ET (100-Gigabit Ethernet) interfaces and OT interfaces:

  - **/junos/system/linecard/optical**
  - **/junos/system/linecard/otn**

  To provision the sensor to export data through gRPC, use the **telemetry Subcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

  [See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **RPM and TWAMP statistics sensor support for Junos Telemetry Interface (JTI) (MX Series)**—Starting with Junos OS Release 18.3R1, you can export Two-Way Active Management Protocol (TWAMP) and real-time performance monitoring (RPM) statistics. TWAMP (described in RFC 5357) and RPM are two methods to measure traffic performance between endpoints. These methods work by sending active probe packets and measuring parameters such as packet loss, delay, and jitter between the endpoints. These statistics provide RPM and TWAMP monitoring data results collected by Juniper devices. You can use this information to ensure service level agreements, improve network design, and optimize traffic engineering.

  Export TWAMP and RPM statistics through JTI using gRPC streaming. The following resource paths are supported:

  - **/junos/rpm/probe-results/probe-test-results/**
  - **/junos/rpm/history-results/history-single-test-results/**
  - **/junos/rpm/server/active-servers/**
  - **/junos/twamp/client/control-connection/**

- **/junos/twamp/client/probe-test-results/**
- **/junos/twamp/client/history-test-result/**
- **/junos/twamp/server/control-connection/**

To provision the sensor to export data through remote procedure call (gRPC) streaming, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

Beginning in Junos OS Release 18.2X75-D10, OpenConfig and Network Agent packages are bundled into the Junos image by default. OpenConfig can be found as a default package named junos-openconfig, and Network Agent content exists in the Junos as a daemon through the na-telemetry package. The OpenConfig package can still be installed as an add-on package on top of the default package if you want to upgrade OpenConfig without upgrading Junos OS.

[See Configuring a Junos Telemetry Interface Sensor (CLI Procedure) and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **Expanded physical interface queue and traffic statistics sensors for Junos Telemetry Interface (JTI) (PTX, MX, EX, QFX, ACX)**—Starting with Junos OS Release 18.3R1, additional resource paths are added to stream physical (IFD) statistics.

  Prior to Junos OS Release 18.3R1, both traffic and queue statistics for physical interfaces (IFD) are sent out together using the resource path **/interfaces** for gRPC streaming (which is internally used to create **/junos/system/linecard/interface/**) or **/junos/system/linecard/interface/** for UDP (native) sensors.

  Now, traffic and queue statistics can be delivered separately. Doing so can reduce the reap time for non-queue data for platforms supporting Virtual Output Queues (VOQ).

  The following UDP resource paths can be configured:

  - **/junos/system/linecard/interface/** is the existing resource path (no change). Traffic and queue statistics are sent together.
  - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
  - **/junos/system/linecard/interface/queue/** exports queue statistics.

  The gRPC resource path **/interfaces** now has the following behavior:

  - In releases prior to Junos OS 18.3R1, it delivers all IFD traffic and queue statistics. In Junos OS 18.3R1 and higher, it delivers statistics in two sensors:
    - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
    - **/junos/system/linecard/interface/queue/** exports queue statistics.

  To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS

module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

[See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

- **Expanded ON_CHANGE support for LLDP telemetry data through Junos Telemetry Interface (JTI) (MX Series)**—Starting with Junos OS Release 18.3R1, OpenConfig support through remote procedure calls (gRPC) and JTI is expanded to support additional ON_CHANGE support for LLDP telemetry sensors. Periodical streaming of OpenConfig operational states and counters collects information at regular intervals. ON_CHANGE support streams operational states as events (only when there is a change), and is preferred over periodic streaming for time-sensitive missions.

  When you create a subscription using a top-level container as the resource path (for example, **/lldp**), a leaf under the resource path **/lldp** with ON_CHANGE support is automatically streamed based on events. If a leaf under the resource path does not have ON_CHANGE support, it will not be streamed.

  These paths, previously supporting periodical streaming only, now also support ON_CHANGE streaming:

  - **/lldp/state/enabled**
  - **/lldp/state/chassis-id**
  - **/lldp/state/chassis-id-type**
  - **/lldp/state/system-name**
  - **/lldp/state/system-description**
  - **/lldp/state/hello-timer**
  - **/lldp/interfaces/interface/state/name**
  - **/lldp/interfaces/interface/state/enabled**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/chassis-id**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/chassis-id-type**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/port-id**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/port-id-type**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/port-description**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/system-name**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/system-description**
  - **/lldp/interfaces/interface/neighbors/neighbor/state/management-address**

- **/lldp/interfaces/interface/neighbors/neighbor/state/management-address-type**
- **/lldp/interfaces/interface/neighbors/neighbor/capabilities**

These resource paths from the preceding list do not change with an event, but will be streamed on creation and deletion:

- **/lldp/interfaces/interface/neighbors/neighbor/state/chassis-id**
- **/lldp/interfaces/interface/neighbors/neighbor/state/chassis-id-type**
- **/lldp/interfaces/interface/neighbors/neighbor/state/system-name**

Before events are streamed, there is an initial stream of states to the collector, followed by an END_OF_INITIAL_SYNC. This notice signals the start of event streaming.

To provision the sensor to export data through gRPC streaming, use the **telemetry Subscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

Beginning in Junos OS Release 18.2X75-D10, OpenConfig and Network Agent packages are bundled into the Junos image by default. OpenConfig can be found as a default package named junos-openconfig, and Network Agent content exists in the Junos as a daemon through the na-telemetry package. The OpenConfig package can still be installed as an add-on package on top of the default package if you want to upgrade OpenConfig without upgrading Junos OS.

[See Configuring a Junos Telemetry Interface Sensor (CLI Procedure) and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **BGP and statically configured SR-TE traffic statistics sensor support for Junos Telemetry Interface (JTI) (MX Series and PTX Series)**—Starting with Junos OS Release 18.3R1, you can export traffic statistics for both ingress IP traffic and transit MPLS traffic that take Segment Routing Traffic Engineering (SR-TE) paths. This feature provides support for BGP [DRAFT-SRTE] and statically configured SR-TE policies at ingress routers.

  Export JTI statistics using either gRPC streaming or UDP native sensors. The following resource paths are supported.

  For UDP native sensors:

  - **/junos/services/segment-routing/traffic-engineering/ingress/usage/**
  - **/junos/services/segment-routing/traffic-engineering/transit/usage/**

  For gRPC streaming:

  - **/mpls/signaling-protocols/segment-routing/**

  For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

To provision the sensor to export data through gRPC streaming, use the **telemetry Subscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

For both export methods, you also must specify that these statistics be collected. To do this, configure collection at the **[edit protocols source-packet-routing telemetry statistics]** hierarchy level.

[See sensor, source-packet-routing, and Guidelines for gRPC Sensors (Junos Telemetry Interface).

MPLS

- **Support of pseudowire headend termination for L3VPN and MVPN (MX Series)**—Starting in Junos OS Release 18.3R1, the support for pseudowire subscriber service interface over redundant logical tunnels is introduced in Layer 3 VPNs and draft-rosen multicast VPNs. Earlier, Layer 3 VPNs provided support for pseudowire subscriber services over logical tunnel interfaces only, and these interfaces used unicast routing protocols, such as OSPF and BGP. This feature introduces provisioning of a multicast routing protocol, Protocol Independent Multicast (PIM), on the pseudowire subscriber interfaces, which gets terminated on the virtual routing and forwarding (VRF) routing instance.

  With this feature, you can enable pseudowire subscriber interfaces for inet, inet6, dual inet, and inet6 stack families, and benefit from the additional resiliency support because of the increase in pseudowire logical interface devices scaling numbers.

  [See Pseudowire Subscriber Logical Interfaces Overview.]

Multicast

- **Persistent designated-router status for last-hop routers (MX Series)**—Starting in Junos OS Release 18.3R1, you can configure a designated router to persist according to your design criteria rather than according to the results of the default designated-router election logic by setting the **stickydr** CLI command.

  Use **stickydr** to prevent traffic loss, for example, in situations where the designated router election may result in unintended changes after an interface-down event or device upgrade.

  To enable designated-router persistence on a configured LAN, enable **stickydr** on all last-hop routers in the LAN, as shown in the following example:

  ```
  set routing-instances instance-name protocols pim interface interface-name stickydr
  ```

  [See stickydr.]

Network Management and Monitoring

- **Customized MIBs for sending custom traps based on syslog events (MX Series)**—Starting in Junos OS Release 18.3R1, there is a process whereby customers can define their own MIBs for trap notifications. The customized MIB maps a particular error message with a custom OID rather than a generic one. Juniper Networks provides two new MIB roots reserved for customer MIBs, one for the custom MIB

modules and the other for the trap notifications. For this process, you must convert the MIB to YANG format, and a tool is available for that.

[See Customized SNMP MIBs for Syslog Traps.]

- **Support over aggregated Ethernet interfaces added for SNMP CoS MIB for interface-sets queue counters (MX Series)**—Starting in Junos OS Release 18.3R1, Junos OS supports SNMP reporting of queue statistics for static interface-sets configured over Aggregate Ethernet (AE) interfaces.

  [See show snmp mib and SNMP MIB Explorer.]

Restoration Procedures Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (MX Series)**—Starting in Junos OS Release 18.3R1, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode. The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

  [See Saving a Rescue Configuration File.]

Routing Protocols

- **IPv4 over IPv6 tunnel scaling per chassis by increasing number of line cards**—Starting in Junos OS Release 18.3R1, you can configure BGP to tunnel the IPv4 unicast routes along with IPv6 nexthop.
- **Junos OS, OpenConfig, and Network Agent packages are delivered in a single TAR file (MX Series)**—Starting in Junos OS Release 18.3R1, the Junos OS image includes the OpenConfig package and Network Agent; therefore, you do not need to install OpenConfig or Network Agent separately on your device.

  [See Installing the OpenConfig Package, and Installing the Agent Network Package.]

- **IS-IS overloading stub networks (MX Series)**—Starting in Junos OS Release 18.3R1, new configuration options **external-prefixes**, and **internal-prefixes** are available at the **[edit protocols isis overload]** hierarchy level to control overload of internal prefixes, external prefixes or both internal and external prefixes according to network requirements. The user can choose not to receive any traffic for internal and external prefixes advertised by the overloaded IS-IS routers unless the router is the only node in the network that hosts the prefix. In previous Junos OS releases, overloaded IS-IS routers continued to receive traffic for prefixes even if the user did not want to receive traffic for directly attached prefixes.

  [See Configuring IS-IS Prefix Overload .]

Security

- **Support for configuring MACsec EAPoL destination address (MX Series)**—Starting in Junos OS Release 18.3R1, you can configure an Extensible Authentication Protocol over LAN (EAPoL) destination MAC

address by including the **eapol-address (pae | provider-bridge | lldp-multicast)** statement at the **[set security macsec connectivity-association *connectivity-association-name* mka]** hierarchy level.

To establish a MACsec session, MACsec Key Agreement PDUs are sent or received between nodes. These PDUs are EAPoL packets and, by default, their destination MAC address is the EAPoL multicast address 01:80:C2:00:00:03. If the nodes are connected through a provider network, they might consume these multicast packets or drop them depending on their configuration. To overcome this issue, configure the EAPoL address for PAE, provider-bridge, and LLDP multicast by using the aforementioned configuration.

[See mka (MX Series).]

- **Support for AES-256 MACsec encryption (MX80, MX104, MX240, MX480, and MX960)**—Starting in Junos OS Release 18.3R1, the MIC-MACSEC-20G MIC provides 256-bit MACsec encryption on MX80, MX104, MX240, MX480, and MX960 routers. This MIC supports MACsec on twenty 1-Gigabit Ethernet SFP ports and on two 10-Gigabit Ethernet SFP+ ports in the following hardware configurations:
  - Installed directly on the MX80 and MX104 routers
  - Installed on MPC1, MPC2, MPC3, MPC2E, MPC3E, MPC2E-NG, and MPC3E-NG line cards on the MX240, MX480, and MX960

Service Applications

- **Support for filtering DNS requests for blacklisted website domains (MX Series with MS-MPCs)**—Starting in Junos OS Release 18.3R1, you can configure DNS filtering to identify DNS requests for blacklisted website domains.

  For DNS request types A, AAAA, MX, CNAME, TXT, SRV, and ANY, you also configure the action to take for a DNS request for a blacklisted domain. You can either:

  - Block access to the website by sending the client a DNS response corresponding to the DNS request type with the IP address or fully qualified domain name (FQDN) of a DNS sinkhole server. This ensures that the client sends further traffic for the blacklisted domain to the sinkhole server.
  - Log the request and allow access.

  For other DNS request types for a blacklisted domain, the request is logged and access is allowed.

  [See Filtering DNS Requests for Blacklisted Website Domains.]

- **MX Series Virtual Chassis NAT support (MX240, MX480, and MX960 routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 18.3R1, you can configure a two-member MX Series Virtual Chassis to use IPv4-to-IPv4 basic network address translation (NAT), dynamic NAT, static destination NAT, dynamic NAT with port mapping, and stateful NAT64. A two-member MX Series Virtual Chassis configuration supports a maximum of four MS-MPCs and four MS-MICs per Virtual Chassis.

  [See Protocols and Applications Supported by the MS-MIC and MS-MPC.]

Software-Defined Networking

- **Support for PCE-initiated point-to-multipoint LSPs (MX Series)**—Starting in Junos OS Release 18.3R1, the Path Computation Element Protocol (PCEP) functionality is extended to allow a stateful PCE to initiate, provision, and modify point-to-multipoint traffic engineering LSPs through a PCC.

  Currently, Junos OS supports only point-to-point PCE-initiated LSPs. With the introduction of point-to-multipoint PCE-initiated LSPs, a PCE can initiate and provision a point-to-multipoint LSP dynamically without the need for local LSP configuration on the PCC. The PCE can also control the timing and sequence of the point-to-multipoint path computations within and across (PCEP) sessions, thereby creating a dynamic network that is centrally controlled and deployed.

  [See Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs.]

- **Support for Junos Node Slicing (MX2008 routers)—**Starting with Junos OS Release 18.3R1, Junos Node Slicing is supported on MX2008 routers. Junos Node Slicing allows a single MX Series router to be partitioned to appear as multiple, independent routers. Each partition has its own Junos OS control plane, which runs as a virtual machine (VM), and a dedicated set of line cards. Each partition is called a guest network function (GNF). In the node slicing setup, the MX Series router functions as the base system (BSYS).

  [See Junos Node Slicing Overview.]

- **Abstracted Fabric Interface Support for Junos Node Slicing (MX Series Routers with MPC5E and MPC6E)**—Junos Node Slicing supports Abstracted Fabric (AF) interface, a pseudo interface that facilitates routing control and management traffic between guest network functions (GNFs) via the switch fabric. An AF interface is created on a GNF to communicate with its peer GNF when the two GNFs are configured to be connected to each other. The bandwidth of the AF interfaces changes based on the insertion or reachability of the remote line card or MPC. Starting in Junos OS Release 18.3R1, GNFs support the following AF-capable MPCs as well:

  - MPC5E (MPC5E-40G10G, MPC5EQ-40G10G, MPC5E-40G100G, MPC5EQ-40G100G)
  - MPC6E (MX2K-MPC6E)

  See [Abstracted Fabric (AF) Interface.]

- **Support for transmit load-balancing statistics on abstracted fabric interface (MX Series)—**Starting in Junos OS Release 18.3R1, Junos Node Slicing supports transmit load-balancing statistics on abstracted fabric (AF) interfaces. The **show interfaces *af-interface-name*** output provides transmit statistics of each Packet Forwarding Engine peer list on a given AF interface, in addition to the physical interface statistics. The output displays information such as residual transmit statistics, fabric queue statistics, and residual fabric queue statistics.

  [See show interfaces (Abstracted Fabric).]

- **Support for non-root users in JDM for Junos Node Slicing—**Starting in Junos OS Release 18.3R1, Juniper Device Manager (JDM) for Junos Node Slicing supports configuration of non-root users. A JDM root user can create non-root users in the JDM by using the **set system login user *username* class *class***

command. The non-root users can interact with JDM; orchestrate and manage the GNFs; and monitor the state of the JDM, host server, and the GNFs by using the existing JDM CLIs.

[See Configuring Non-Root Users in JDM (Junos Node Slicing).]

- **Support for OpenDaylight controller (Nitrogen) (MX Series)** —Starting with Junos OS Release 18.3R1, MX Series routers support the Nitrogen release version of the OpenDaylight (ODL) controller. The ODL controller, also known as ODL platform, provides a southbound Network Configuration Protocol (NETCONF) connector API, which uses NETCONF and YANG models to interact with a network device. You can use the ODL controller to carry out configuration changes in MX Series routers, and provision and orchestrate the routers. The ODL controller provides an open-source platform for network programmability aimed at enhancing software-defined networking (SDN).

  [See Configuring Interoperability Between MX Series Routers and OpenDaylight.]

Subscriber Management and Services

- **DHCPv6 subscriber class differentiation with the DHCPv6-Options VSA (26-207) (MX Series)**—Starting in Junos OS Release 18.3R1, you can use VSA 26-207 to differentiate between different classes of subscribers during DHCPv6 relay authentication. Configure your RADIUS server to include the following information in DHCPv6 Option 17:

  - Juniper Networks enterprise number, 2636
  - Suboption 5, JDHCPD_VS_OPT_CODE_KT_SUBSCRIBER_CLASS

  You set a different value for suboption 5 for each class. The VSA conveys this Option 17 information in the Access-Accept message RADIUS returns during DHCPv6 subscriber authentication. The DHCPv6 relay agent extracts the Option 17 information and passes the information to the DHCPv6 local server in the Relay-Forward header.

  In earlier releases, only the DHCP local server supports VSA 26-207; only suboption 1 (hostname) and suboption 4 (location) are supported.

  [See Exchange of DHCPv4 and DHCPv6 Parameters with the RADIUS Server Overview.]

- **Support for per-subscriber application-aware policy control (MX Series routers with Junos Node Slicing)**—Starting in Junos OS Release 18.3R1, Junos Node Slicing supports per-subscriber application-aware policy control. With this support, the Multiservices MPCs and Multiservices MICs on the routers configured with Junos Node Slicing provide per-subscriber policy control based on Layer 7 application identification information for the IP flow (for example, YouTube) or Layer 3 and Layer 4 information for the IP flow (for example, the source and destination IP address). Subscriber application-aware policy actions can include:

  - Redirecting HTTP traffic to another URL or IP address
  - Steering with a routing instance
  - Setting the forwarding class
  - Setting the maximum bit rate

- Setting the gating status to blocked or allowed
- Setting the allowed burst size
- Logging data for subscriber application-aware data sessions and sending that data in an IP Flow Information Export (IPFIX) format to an external log collector, using UDP-based transport.

[See Understanding Application-Aware Policy Control for Subscriber Management.]

- **Support for remote device service management (MX Series)**—Starting in Junos OS Release 18.3R1, a new service type is supported on BNGs, remote-device-services. The new remote device services manager (RDSM) provisions and deprovisions services on remote devices that are managed as logical extensions to the BNG. Remote devices are subscriber-facing devices such as OLTs, DSLAMs, and other access nodes. The BNG acts as a proxy server for the remote devices for service configuration. To external management and provisioning (PCRF, RADIUS) systems, the BNG together with its remote devices acts as a single addressable network element. A dynamic service profile is applied by the external authority by reference during subscriber provisioning to initiate service actions on the remote devices.

  [See Remote Device Services Manager (RDSM) Overview.]

- **Enhancements to static subscriber usernames and interface support (MX Series)**—Starting in Junos OS Release 18.3R1, the following enhancements are added for subscribers on static interfaces:

  - You can include outer and inner VLAN tags from the static interface in the global or group usernames.
  - You can specify any single character as the delimiter between username elements.
  - Pseudowire interfaces over logical tunnels are supported for static subscribers, which enables full subscriber management equivalent to dynamic subscribers for statically provisioned subscribers whose traffic is transported over IP/MPLS access models (PS/LT).

    The maximum logical unit number range for pseudowire static interfaces is increased from 16,385 to 1,073,741,823.

  [See Configuring the Static Subscriber Global Username and Configuring the Static Subscriber Group Username.]

- **Support for IPFIX mediation on the BNG (MX Series)**—Starting in Junos OS Release 18.3R1, you can configure a BNG to act as an intermediary device between IPFIX exporters and collectors, while having the functions of both. The IPFIX mediator function collects performance management data via IPFIX records from downstream access network devices such as OLTs and advanced ONUs. This data along with local performance management data from the MX BNG is aggregated and relayed to an upstream IPFIX collector. From the reference point of the IPFIX collector, IPFIX mediation enables the router and its associated access network devices to appear as a single IPFIX export source leveraging a single TCP/IP connection between the MX BNG and the upstream collector.

  [See IPFIX Mediation on the BNG.]

- **Support for TCP port forwarding (MX Series)**—Starting in Junos OS Release 18.3R1, TCP port forwarding enables a BNG to mediate communication between its connected access nodes and service provider back-office systems, such as external management and provisioning systems (leveraging NETCONF XML

management protocol) and TACACS+ servers. The BNG and its downstream access nodes are presented to back-office systems as a single addressable network element. Communication requests to and from access nodes are redirected from one address and port number combination to another while packets traverse the MX Series router. You configure unique combinations of listening ports and listening addresses on the BNG. TCP connections are triggered when traffic from acceptable prefixes arrives on the listening port and matching listening address.

[See TCP Port Forwarding for Remote Device Management.]

SEE ALSO

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for MX Series routers.

Class of Service (CoS)

- **Junos commit notification of unsupported configuration**—Junos OS does not support changing the **hierarchical-scheduler** mode of a logical tunnel interface, or redundant logical tunnel interface, if an active pseudowire subscriber interface is attached to it. A commit error has now been added to provide the notification.

EVPN

- **Support for an VNI of zero**—Starting with Junos OS Release 18.3R2, Junos supports using a VXLAN Network Identifier (VNI)=0 when configuring a bridge domain or vlan in an EVPN-VXLAN network.
- **Changes in encoding the ESI label field (MX Series)**—Starting in 18.3R2, Junos OS switched from using lower-order bits to higher-order bits in encoding the ESI label field. This results in BUM traffic loss and duplication in traffic. If you encounter this, and you wish to use a mix of Junos OS releases, you must include the **es-label-oldstyle** statement at the **[edit routing-instances *routing-instance-name* protocols evpn]** hierarchy on the device that is running the Junos OS release that supports higher-order bit encoding of the ESI label.

General Routing

- **Zero MAC address (00:00:00:00:00:00) treated as "my mac" (MX-Series)**—When an Ethernet packet arrives in ingress, pre-classifier engine will perform a lookup of MAC address. If the MAC address matches an entry in the pre-classifier Ternary Content Addressable Memory (TCAM) and the entry has “my mac” attribute, pre-classifier engine will set the “my mac” bit in the cookie prepended to the incoming packet. In current implementation, MAC address “00:00:00:00:00:00” (zero MAC) is programmed as default value for “my mac” TCAM entries when the pre-allocated entries are not used or configured. Hence the packets with zero MAC are marked as “my mac” in the packet cookie. Forwarding engine will check “my mac” bit in the packet cookie. If “my mac” bit is 0, the packet will be dropped. If “my mac” bit is 1, further L2, L3, MPLS lookup will be performed. The “my mac” behavior is applicable since the day one release.
- **Root XML tag change for show rsvp pop-and-forward | display xml command (MX480)**—We’ve changed the root XML tag for the show rsvp pop-and-forward | display xml command to

rsvp-pop-and-fwd-information to make it consistent with the XML tag convention. In earlier releases, the command output displays rsvp-pop-and-fwd-info XML tag. Update the scripts with the rsvp-pop-and-fwd-info XML tag to reflect the new rsvp-pop-and-fwd-information XML tag.

[See Junos XML API Explorer - Operational Tags.]

Interfaces and Chassis

- **Error thrown when router configuration updated on live system**—In Junos OS Release 18.3R1, on MX Series routers with the RE-S-X6-64G and RE-MX2K-X8-64G Routing Engines, when the user changes the router configuration on a live system, or when the user deletes an interface that has active traffic, the message **select: protocol failure in circuit setup** is randomly displayed. However, there is no known functional impact.
- In MX204 routers, the error messages are logged when **vlan-tagging** for a trunk interface that is not configured. These error messages were previously logged with severity level “critical” even though they were not critical enough to require immediate action. The maximum transmission unit (MTU) of interface with or without VLAN-tagging is now logged in as the informational error message (instead of critical error message).
- **IRB not supported on Pseudowire Subscriber (PS) Logical Interface in bridge-domain (MX Series)**—In Junos OS Releases 18.3R2, integrated routing and bridging (IRB) is not supported on Pseudowire Subscriber (PS) logical Interface. Hence you cannot add IRB to bridge domain with PS interface, that is, you cannot configure IRB and PS interface in the same bridge domain.

  Note that adding IRB to a bridge-domain having Pseudowire Subscriber (PS) Logical Interface causes kernel crash and continuous reboot of the router until the configuration is rolled back.

  > **NOTE:** IRB is not supported on PS only in bridge-domain.

  [See bridge-domain.]
- **Support for MAP-E encapsulation and decapsulation on Inline Service Interfaces (MX2010)**—In Junos OS Releases 18.2R3 and 18.3R2 the MX2010 routers support encapsulation and decapsulation of the following ICMP message types for inline service (si) interfaces:
  - Time Exceeded (type 11)
  - Destination unreachable (type 3)
  - Source quench (type 4)
  - Parameter problem (type 12)

- Address mask request and Address mask reply (type 17 and type 18)
- Redirect (type 5)

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (MX Series)**—In Junos OS Release 18.3R2, the **show lacp interfaces | display xml** command displays a new XML tag element **<lacp-hold-up-state>**. The **<lacp-hold-up-state>** displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces were in a single **<lacp-hold-up-information>** XML tag. Now, for each interface it is displayed in a separate **<lacp-hold-up-information>** XML tag.

Junos OS XML, API, and Scripting

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (MX Series)**—Starting in Junos OS Release 18.3R1, Junos OS does not support configuring an MD5 or SHA-1 checksum hash to verify the integrity of local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) scripts or support using an MD5 or SHA-1 checksum hash with the **op url *url* key** option to verify the integrity of remote op scripts.

MPLS

- When the **no-interface-hello** statement is configured under the **[edit protocols rsvp]** hierarchy, and there is no interface-specific configuration for the hello interval, the **show rsvp interface detail** command output displayed the default **HelloInterval** of 9 seconds.

  Starting in Junos OS Release 18.3R1, with a similar configuration, the **HelloInterval** output field displays 0 as the hello interval.

- **Change in get-pm-mpls-lsp-information tag**—Starting in Junos OS Release 18.3R1, the **show performance-monitoring mpls lsp** command output in the YANG module is changed to match the root XML tag for **get-pm-mpls-lsp-information**. The tag change is from **performance-monitor-mpls-lsp-information** to **pm-information**.

- **Change in get-egress-protection-information tag**—Starting in Junos OS Release 18.3R1, the **show mpls egress-protection** command output in the YANG module is changed to match the root XML tag for **get-egress-protection-information**. The tag change is from **egress-protection-information** to **ep-operational-information**.

- **Bandwidth allocation**—For a label-swtiched path (LSP) that has both **bandwidth** and **minimum-bandwidth** for autobandwidth configured under the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level, the LSP bandwidth is adjusted differently.

  The LSP is initiated with the bandwidth value configured under the **bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name]** hierarchy level. At the expiry of the **adjust-interval** timer, the LSP bandwidth gets adjusted based on the traffic flow.

If the bandwidth to be signaled is less than the value configured under the **minimum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the minimum bandwidth.

If the bandwidth to be signaled is greater than the value configured under the **maximum-bandwidth** statement at the **[edit protocols mpls label-switched-path lsp-name autobandwidth]** hierarchy level, then the LSP is signaled only using the maximum bandwidth.

- **Change in command syntax**—Starting in Junos OS Release 18.3R1, the **show ldp database label-requests** command name is changed to **show ldp database-label-requests** with no change to command functionality.
- Previously, when you configured zero (*0*) as the bandwidth of an RSVP interface, the bandwidth value was overwritten with the default interface bandwidth (raw hardware bandwidth), leading to unexpected behavior in the LSP setup. Starting with Junos OS Release 18.3R2, when you configure zero as the bandwidth, **0** is applied as the RSVP bandwidth.

  [See bandwidth (Protocols RSVP).]
- **Loss of traffic over bypass MPLS LSPs**—If RSVP link or node protection is enabled along with global RSVP authentication, there is loss of traffic over bypass MPLS LSPs at the time of local repair, when the point of local repair (PLR) and the merge point devices have different versions of the Junos OS software installed on them. That is, one device is running a release prior to Junos OS Release 16.1, and the other device is running a release starting with Junos OS Release 16.1R4-S12.

Network Management and Monitoring

- **Junos OS does not support management of YANG packages in configuration mode (MX Series)**—Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages using the **run** command in configuration mode is not supported.
- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (MX Series)**—Starting in Junos OS Release 18.3R2, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.
- **Change in severity level of XQSS errors (MX Series)**—Starting in Junos OS Release 18.3R2, on MX series routers with the MPC7E-10G, MPC7E-MRATE, MPC8E, and MPC9E line cards, the severity level of the following errors have been changed from Fatal to Major.
  - XQSS_CMERROR_CPQW_ERR_INT_FSET_SLOW_DEQ_DRY_ERR
  - XQSS_CMERROR_CPQW_ERR_INT_FSET_FAST_DEQ_DRY_ERR

With this change, the above errors no more cause the entire FPC to go offline by default. Instead, these errors cause the affected Packet Forwarding Engine (PFE) to be disabled, as **disable-pfe** is the default action associated with Major errors on MX series routers.

Additionally, the severity level of the correctable error **XQSS_CMERROR_CORRECTABLE_MEM_ERR** has been changed from Fatal to Minor.

You can use the commands **show chassis errors active detail fpc-slot *slot*** and **show chassis fpc errors *slot*** to view more details of, and the default actions associated with, these errors.

[See show chassis fpc errors.]

Routing Protocols

- **IS-IS adjacency SID routes retained only when backup path is available**—Starting in Junos OS Release 18.3R1, when an IS-IS link flaps the adjacency SID routes are retained in the RIB, (also known as the routing table) and the FIB, (also known as the forwarding table) only if a backup path is available. In earlier Junos OS releases, adjacency SID routes were retained in the RIB and FIB even when a backup path was not available.

Security

- **Syslog updated when configuring XPN cipher suite on a non-xpn supported interface (MX Series Routers)**—In Junos OS Release 18.3R1, on MX Series Routers, if you attempt to configure XPN cipher suite (gcm-aes-xpn-128 or gcm-aes-xpn-256) for a connectivity association and attach the connectivity association to an interface on the PIC that does not support XPN cipher suite, then during runtime, a syslog is logged as below (and default non-xpn cipher suite is used):

  macsec_ciphersuite_is_supported MACSec: ifd *ifd_id (ifd_name)*, Cipher suite *cipher id (cipher name)* NOT SUPPORTED.

Services Applications

- **Change in error message displayed while fragmenting or de-fragmenting IPv6 GRE tunnel interface (MX Series routers)**—In Junos OS Release 18.3R2, on a IPv6 GRE tunnel interface, when you enable fragmentation using the **allow-fragmentation** command or disable fragmentation using the **do-not-fragment** command, the following error message is displayed:

  Fragmentation for V6 tunnels is not supported

  In earlier Junos OS releases, the following message was displayed:

  dcd_config_ifl_tunnel:Fragmentation for V6 tunnels is notsupported

- **Support for host generated traffic on a GRE over GRE tunnel (MX Series)**—In Junos OS Release 18.3R2, you can send host generated traffic on a GRE over GRE tunnel. However, when path maximum transmission unit (PMTU) is updated for the outer GRE tunnel, MTU for inner GRE tunnel is not corrected.

Software Installation and Upgrade

- **ZTP is supported on MX PPC platforms (MX Series)**—As of Junos OS Release 18.3R1, zero touch provisioning (ZTP) is supported on MX PPC platforms (which are MX5, MX10, MX40, MX80, and MX104 routers). Before the fix, the ZTP process did not start to load image and configuration for MX PPC routers.

  [See Junos OS Installation Package Names.]

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (MX Series)**—Starting in Junos OS Release 18.3R1, the jdhcpd process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
  - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
  - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

  The new behavior is as follows:

  1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
  2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

  In earlier releases:

  1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
  2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

  [See Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview.]

- **Disabling a pseudowire underlying interface (MX Series)**—Starting in Junos OS Release 18.3R1, you cannot disable the underlying logical tunnel (lt) interface or redundant logical tunnel (rlt) interface when

a pseudowire is anchored on that interface. If you want to disable the underlying interface, you must first deactivate the pseudowire.

[See Configuring a Pseudowire Subscriber Logical Interface Device.]

- **Bandwidth options match for inline services and tunnel services (MX Series)**—Starting in Junos OS Release 18.3R1, you can configure the same bandwidth options for inline services with the **bandwidth** statement at the **[edit chassis fpc *slot-number* pic *number* inline-services** hierarchy level as you can configure for tunnel services with the **bandwidth** statement at the **[edit chassis fpc *slot-number* pic *number* tunnel-services]** hierarchy level.

  [See bandwidth (Inline Services) and bandwidth (Tunnel Services)]

- **ICMP error message rate limit increased (MX Series)**—Starting in Junos OS Release 18.3R2, the maximum rate limit for generating ICMP messages for IPv4 and IPv6 packet errors is increased from 50 pps to 1000 pps. The rate limit applies only to non-ttl-expired packets.

- **Subscribers allowed to log in with bad framed route (MX Series)**—Starting in Junos OS Release 18.3R2, users are allowed to log in if the framed route received from RADIUS is bad; for example, if the format is incorrect. In earlier releases, the subscriber is not allowed to log in. For customers that use multiple framed routes, the new behavior enables the subscriber to have partial access to the network using the routes that are accepted instead of not being allowed any access.

- **Out-of-address SNMP trap requires thresholds to be configured (MX Series)**—Starting in Junos OS Release 18.3R2, the behavior has changed for generating an out-of-address SNMP trap for an address pool configured at the **[edit access address-assignment]** or **[edit routing-instance *name* address-assignment]** hierarchy levels. You must now configure both the high-utilization and abated-utilization thresholds. When the number of assigned addresses surpasses the high-utilization threshold, a high-utilization trap is generated. If all the addresses are assigned from the pool, an out-of-address trap is generated and an out-of-address syslog message is sent.

  In earlier releases, an out-of-address trap is generated when the address pool is exhausted, regardless of whether the thresholds are configured.

  If the number of assigned addresses subsequently drops below the abated-utilization threshold, an abate-high-utilization trap is generated; this behavior is unchanged.

VPNs

- **Output of show l2vpn connections extensive command in XML**—Starting in Junos OS Release 18.3R1, the output for **show l2vpn connections extensive | display xml** will correctly display the output in XML.

SEE ALSO

New and Changed Features | 98

Known Behavior

IN THIS SECTION

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 for MX Series routers.

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

- There is an interface for a direct route starting in ifdown condition. The remote side is then brought up, so I/F goes to ifup. Since it is a direct route, rpd does not install the route or nexthop. It receives that information from the kernel, and just updates a nexthop in rpd local storage. route and nexthop for the interface are taken care of in the kernel. There is no route change in rpd. route_record depends on route flash to find out about updates. Since there is no route change, there is no route flash, so route_record is blissfully unaware. In order to change this, we would need to decide that we want a route flash for this case. Currently, for direct and local routes / nexthops, these are "don't care" in rpd, as far as route updates go. We just update our nexthop info, without marking for any other notifications. A complication for the solution is a change that was done for PR 1002287, where if the NOTINSTALL flag is set, do not send the update to srrd. That flag is set for direct and local routes. Incidentally, this is day-one operation. If the interface is up at startup, it should work correctly. FIB table can provide OIF/GW only. SRC_MASK, DST_MASK, SRC_AS and DST_AS are not available in PFE FIB Table. So SRRD connection is required. Listening to both SRRD and FIB table, and consolidating information will complicate implementation. Scanning entire FIB Table just for the few such routes will have performance impact and will complicate present implementation. This is day 1 implementation for SRRD/Sampled. There are two possible workarounds:

  1) Have the far end interface up when the DUT interface is brought up. In the case where that is not happening, a recovery would be to disable the DUT interface, then enable it again. At that point, everything should be initially brought up in the state we are looking for.

  2) Enable the **nexthop-learning** command. Please refer to the documentation for information on this command. PR1224105

High Availability and Resiliency

- MX Virtual Chassis configurations cannot use ISSU to upgrade from Junos OS Release 18.2R1 to 18.3R1.
- If a MX10003 router is plugged in with SFP/SFP+ via a QSA adapter, then ISSU will not function.
- MX series routers with MPC7E MPCs installed cannot use ISSU to upgrade to Junos OS Release 18.3R1.

General Routing

- When some route or next hop has been created by the application, it is assumed that it can propagate to the rest of the system. KRT asynchronously picks up this state for propagation. There is no reverse indication to the application, if there was an error in propagating the state. The system is supposed to eventually reconcile. So, if SPRING-TE produces a <route, NH> pair that looks legal from the app standpoint, but KRT is not able to download it to the kernel, because the kernel rejected the next hop, the <route, NH> gets stuck in rpd. In the meantime, the previous version of the route (L-ISIS in this case) that was downloaded still lingers in the kernel and Packet Forwarding Engine. PR1253778
- CFM is not supported for L2-over-GRE tunnel. CCM can pass through as transit traffic through GRE interfaces transparently using data path. Link trace functionality uses MAC-learning and re-injecting LTM on GRE interface in case the bridge is configured with CFM. This is not a supported feature. PR1275833
- On MX104 **JTASK_SCHED_SLIP** is seen on commit randomly. PR1281016
- Support for enterprise profile is only provided for 10-Gigabit Ethernet interfaces. Use of 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. PR1310048
- The input int field of the MPLS version 4 data records reports the SNMP index value of the LSI interface instead of the ingress physical interface. PR1312047
- When cmerror disables Packet Forwarding Engine, it does not power off the EA and HMC chips. The periodic continues monitoring the temperature on HMC and other devices. If the temperature is overheated, the system can take proper actions, such as increase the fan speed or shut down the systems. The periodic calls hmc_eri_config_access() to get temperature readings. It is expected to get ERI timeout continuously in this case. PR1324070
- Hardware watchdog does not work on QFX10008 and QFX10002-60C/PTX10002-60C. PR1343131
- After disabling the laser for CWDM optics, optics diagnostics will not report o/p power low and laser current low alarm/warnings. PR1349258
- On a PTX1000 router, after the system is rebooted by issuing the **request vmhost reboot** command, the netproxy service might fail to start. PR1365664
- FPC bounce is required for a mode change from 1-Gigabit to 10-Gigabit PIC speed or vice versa to take effect on the MIC-MACSEC-20G. PR1373400
- Port-level speed configuration is not supported for 10 Gbps mode on the 2x 10Gigabit Ethernet SFPP / 20x Gigabit Ethernet SFP MACsec MIC: Only the pic-mode configuration under the **set chassis fpc <x> pic <a> pic-mode** is to be used to set the PICs in 10-Gigabit speed. PR1373473
- The MIC-MACSEC-20G supports 10 G speed via the **set chassis fpc x pic y pic-mode 10G** configuration applied to both the PICs in that MIC. Any other PIC mode configuration should be removed and then the 10G PIC mode configuration is to be applied. PR1374680
- The 10 Gbps speed-capable ports of the MIC-MACSEC-20G MIC might show the link status as up while the peer side might remain down. PR1382024

- IDS aggregate configuration knob will not be considered for the installation of the IDS dynamic filter. PR1395316
- Junos OS has a limitation that the ARP/NDP state and the associated kernel routes (destination routes) will not be cleared if the ARP/NDP-created next hop has references from RPD. This might impair the clearing of ARP/NDP state via **clear** commands or interface down when host routes are added to the FIB. As a result, it is recommended that a FIB policy be configured to reject host routes before enabling host-route-generation. PR1415400
- Rpd maintains nexthops in a database. Routes from different protocols point to these nexthops. When a route needs to be added to forwarding-table, RPD installs the nexthop and gets a nexthop-index for the installed next hop. The route is installed to the forwarding-table with this next-hop index. Note that in RPD next-hop installation to forwarding-table is need-based that is, when the first route using the nexthop needs to be installed to the forwarding-table, the nexthop is installed. Each nexthop maintains a counter (reference-count) to track its usage by various applications. A nexthop that was installed in the forwarding-table is deleted when there are no users for this next hop, in other words, reference-counts reaches 0. Only when reference-count of a nexthop reaches 0, the nexthop will be deleted from kernel. Due to this logic, there might be a situation where a nexthop remains in forwarding-table even if there are no users of that nexthop in the forwarding-table. For example, A nexthop NH1 was installed in the forwarding-table as part of installation of route R1. Later route R1 becomes inactive in RIB or is blocked by the forwarding policy. However, route R1 still exists in the routing table and still points to the nexthop NH1 that is, there are still users of next hop NH1 in rpd and is still referred. So this nexthop does not get deleted from the forwarding table. PR1415935

Interfaces and Chassis

- Previously, the same IP address could be configured on different logical interfaces from different physical interfaces in the same routing instance (including the master routing instance), but only one logical interface is assigned with the identical address after commit. There is no warning during the commit, only syslog messages indicating an incorrect configuration. PR1221993
- If you configure 64K bridge-domains, with each BD having 2 IFLs and 1 irb interface, you might run into heap memory exhaustion as this requires more than the supported memory on the FPC. As a workaround, configure interfaces in the trunk mode that allow all 4000 vlans, reducing the need to configure IFLs for each BD. The trunk ports are configured in the default instance or for each routing-instance. PR1348363
- At JDM install time, each JDM instance generates pseudo random MAC addresses to be used for JDM's own management interface and for the associated GNFs' management interfaces. At GNF creation time, each GNF instance generates pseudo random MAC addresses to be used as the chassis MAC address pool for the forwarding interfaces of that GNF. Once generated, JDM and GNF MAC addresses are persistent, and will only be deleted when the JDM or GNF instance itself is deleted.

  At a GNF, the Junos OS CLI command **show chassis mac-addresses** can be used to examine its chassis MAC address pool, and the Junos OS CLI command **show interfaces fxp0** can be used to examine the MAC address of its management interface.

At JDM, the CLI command **show interfaces jmgmt0** can be used to examine the MAC address of its management interface.

In case of MAC address duplication across JDM or GNF instances, you must delete and then reinstall the respective JDM or GNF instance and check again for duplication.

- **Error thrown when router configuration updated on live system**—In Junos OS Release 18.3R1 and 18.3R2, on MX Series routers with the RE-S-X6-64G and RE-MX2K-X8-64G Routing Engines, when the user changes the router configuration on a live system, or when the user deletes an interface that has active traffic, the message **select: protocol failure in circuit setup** is randomly displayed. However, there is no known functional impact.
- On MX Series routers, the **request support information** command executes the following show commands in addition to the existing show commands:
  - **show chassis fabric summary**
  - **show chassis fabric fpcs**
  - **show chassis fabric plane**
  - **show chassis fabric reachability**
  - **show chassis fabric degradation**
  - **show chassis fabric destinations**
  - **show chassis fpc**
  - **show chassis power**
  - **show pfe statistics traffic**
- The two SFP+ ports on the Routing Control Board (RCB) of an MX2008 router have two port LEDs each - one Link Status LED and one Link Activity LED per port. On an MX2008 router, which is connected to an external x86 server in a Junos Node Slicing setup, behavior of these LEDs with regard to Junos Node Slicing configuration is as follows:
  - The Link Status LEDs and Link Activity LEDs on both the ports are off when Junos Node Slicing is disabled or not configured.
  - When you have configured **network-slices** on the router (also called base system or BSYS) but have not configured guest network functions (GNFs) on the server, the Link Status LED on each port turns green (steady-glow). In this case, the Link Activity LED on each port is off.
  - When you have configured Junos Node Slicing (including GNFs), the Link Activity LED on each port is amber (blinking), while the Link Status LED on each port remains green (steady-glow).

Platform and Infrastructure

- On all Junos platforms, execution of Python scripts through enhanced automation does not work on veriexec images. PR1334425
- It is expected to see few transient FI cell underflow errors during a unified ISSU as long as they do not persist. PR1353904

Port Security

- **MACsec pre-shared CAK cannot be zeros**—In previous releases of Junos OS, it was possible to have an all-zero pre-shared static connectivity association key (CAK). In this and future releases of Junos OS, all-zero pre-shared CAKs are not allowed. Manually entered all-zero configured keys will not commit, and any such inherited configurations will be automatically nulled during system upgrade.

  Pre-shared keys are exchanged between two devices at each end of a point-to-point link to initiate the MACsec Key Agreement (MKA) protocol and enable MACsec using static CAK security mode. The exact Junos statement affected is **security macsec connectivity-association <name> pre-shared-key cak <number>**.

Routing Protocols

- BGP peer flap is seen when Routing Engine switchover is triggered from the old backup Routing Engine. This issue is seen only with higher scales. The issue is related to slow draining out of the new backup socket. PR1325804
- When 32,000 SR-TE policies are configured at once, during configuration time there might be scheduler slips. PR1339829
- The mcsnoopd error messages are seen in logs while adding or deleting IGMP PIM configuration. These are debug messages and are not harmful. PR1371662

Services Applications

- We do not recommend to configure the ms- interface when AMS bundle in one-to-one mode has the same member interface. PR1209660
- Broadband-edge platforms do not support service-set integration with dynamic profiles when the service set is representing a carrier-grade NAT configuration. As a workaround, you can use next-hop service set configurations and routing options to steer traffic to a multiservices interface (ms) interface where NAT functionality can be exercised. The following configuration snippet shows the basics of statically configuring the multiservices interface next hop and a next-hop service set. Traffic on which the service is applied is forced to the interface inside the network by configuring that interface as the next hop. This configuration does not show other routing-options or NAT configurations relevant to your network.

```
routing-options {
    static {
        route 0.0.0.0/0 {
            next-hop ms-3/0/0.1;
            preference 0;
        }
    }
    ...
}
services {
    service-set CGN {
        nat-rules CGN_SAMPLE;
        next-hop-service {
            inside-service-interface ms-3/0/0.1;
            outside-service-interface ms-3/0/0.2;
        }
    }
    nat {
        ...
    }
}
```

[See Configuring Service Sets to be Applied to Services Interfaces.]

Software Defined Networking

- **JDM restart error**—In some cases, restarting Juniper Device Manager (JDM) results in the following error message: **Restarting JDM Job for jdm.service failed because the control process exited with error code. See "systemctl status jdm.service" and "journalctl -xe" for details...** However, JDM automatically recovers from the error condition and restarts successfully. A possible reason for this message is that the control process exited the last session because of an error. In the case of such errors, you can check the operational state of JDM by using the **jdm status** command.
- Starting in Junos OS Release 18.3R2, in Junos Node Slicing, memory allocation to GNFs is shown in gibibytes (GiB), instead of gigabyte (GB). The unit GiB represents memory allocations in multiples of 1024 bytes. This change is applicable only to the JDM CLI help strings at the **[edit virtual-network-functions *vnf-name* resource-template]** hierarchy and to the output of the JDM show command **show virtual-network-functions *vnf-name***.

Subscriber Management and Services

- Before you make any changes to the underlying interface for a demux0 interface, you must ensure that no subscribers are currently present on that underlying interface. If any subscribers are present, you must remove them before you make changes.
- For dual-stacked clients over the same PPP over L2TP LNS session, enhanced subscriber management does not support configurations where both of the following are true:
  - The CPE sends separate DHCPv6 solicit messages for the IA_NA and the IA_PD.
  - The solicit messages specify a type 2 or type 3 DUID (link-layer address).

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

SEE ALSO

New and Changed Features | **98**

Changes in Behavior and Syntax | **116**

Known Issues | **131**

Resolved Issues | **148**

Documentation Updates | **181**

Migration, Upgrade, and Downgrade Instructions | **182**

Product Compatibility | **189**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 for MX Series routers.

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

EVPN

- The Layer 2 address learning process (l2-ald) might generate a core file in a scaled L2 setup, including a bridge domain, VPLS, EVPN, and so on. The l2-ald core file usually follows a kernel page fault that recovers on its own. In some cases, a manual restart of the process is needed to recover the logs: /kernel: %KERN-3-BAD_PAGE_FAULT: pid 69719 (l2ald), uid 0: pc 0x88beb5ce got a read fault at 0x6ca, x86 fault flags = 0x4 /kernel: %KERN-6: pid 69719 (l2ald), uid 0: exited on signal 11 (core dumped) init: %AUTH-3: l2-learning (PID 69719) terminated by signal number 11. A core file is generated. PR1142719
- A core link flap might result in an inconsistent global MAC count. PR1328956
- In an EVPN-VXLAN deployment, the rpd process might crash on the new master Routing Engine after performing a GRES. PR1333754
- In platforms running Junos OS, the l2ald daemon might crash during MAC address processing. The MAC learning process will be impacted during the 12ald crash. The l2ald recovers itself. PR1347606
- Bidirection L2 traffic floods for around 5 seconds for streams from SH to MH, when the **clear mac table** command is executed on the MX Series router because MAC addressing takes time to develop in the system. The **clear mac table** command is a disruptive command that deletes all dynamic MAC addresses in the system. PR1360348
- Type 2 EVPN routes are missing after the EVPN protocol is activated or/and deactivated. PR1362598
- An error is observed when you execute the vxlan ping overlay rpc command. It works through the CLI but not through RPC. PR1373025

- When EVPN is configured with class-of-service-based forwarding (CBF), traffic might be lost for the CBF services. PR1374211
- In an EVPN scenario, even if an IPv6 **virtual-gateway-address** is configured on "IRB" interface, the router advertisement (RA) packets may be sent out with the physical interface/link-local IPv6 address instead of configured virtual-gateway-address. PR1384574
- If the parameter **auto** is set to the statement **vrf-target** within an instance-type of EVPN/virtual-switch, the rpd process would crash after deactivating the autonomous-system (AS) configured. PR1381940
- On activating and deactivating EVPN VXLAN on a routing-instance, the Layer 2 interface on which the MAC is installed is deleted and the MAC entry is lost. PR1387247
- Upon activating nonstop-routing backup routing-engine routing protocol daemon asserting when processing EVPN ESI changed with dynamic-list-next-hop. PR1425687

Forwarding and Sampling

- SRRD daemon acts as a server for all J-Flow clients. The JFlow clients can be either Packet Forwarding Engines or PICs performing JFlow. The maximum number of JFlow clients were previously 32 and it has been increased to 64 in this release. PR1261783
- Heap memory leaks occur on the DPC when the flow specification route is changed. PR1305977
- When customers are trying to configure **flex-match-range** functionality, they receive errors with configuration settings when using bit-length 128 for IPv6. PR1389103
- On Junos Fusion, ingress policing on SD is broken and the configuration statement **set interfaces layer2-policer input-policer <policer-name>** is not supported. PR1395217

General Routing

- ALG-SIP64: SIP session fails when the IPv4 SIP client in a public network initiates a SIP call with the IPv6 SIP client in a private network. PR1139008
- The jl2tpd might generate core files when issuing the CLI command **show services l2tp tunnel statistics**. PR1146771
- NAT64: Source-prefix filtering and protocol filtering of the CGNAT sessions are incorrect. For example, **show services sessions extensive protocol udp source-prefix <0:7000::2>** displays incorrect filtering of the sessions. PR1179922
- GUMEM errors for the same address might continually be logged if a parity error occurs in a locked location in GUMEM. Since GUMEM utilizes ECC memory, any error is self-correcting and has no impact on the operation of the router. In a rare case, such a parity error might appear repeatedly at a specific location. As a workaround, the error can be cleared by rebooting the FPC. PR1200503

- SMID daemon stops responding to management requests after a jl2tpd (L2TP daemon) crash on an MX960 BNG. PR1205546
- Various common situations lead to different views of forwarding information between kernel and Packet Forwarding Engines. For example, **fpc7 KERNEL/PFE APP=NH OUT OF SYNC: "error code 3 REASON: NH add received for an ifl that does not exist ERROR-SPECIFIC INFO: nh_id=562 , type = Hold, ifl index 334 does not exist TYPE-SPECIFIC INFO: none**. No service impact is seen on MPC2 and MPC3 type cards. PR1205593
- In a rare race condition, multiple interrupts are not handled properly on an MX platform with MPC7E, MPC8E, or MPC9E and a PTX platform with FPC3-PTX-U2 or FPC3-PTX-U3, which could lead to core files. As a workaround, the interrupt code is optimized to avoid the unnecessary call to prevent the issue. PR1208536
- When the MPC is removed, the Link Error column in the **show chassis fabric summary extended** output shows YES for all fabric planes. Whereas, when the MPC is taken offline using the CLI command, the output shows correctly. PR1214611
- Load balancing is uneven across aggregate Ethernet (AE) member links when the AE bundle is part of an equal-cost multipath (ECMP) path. The AE member links need to span Virtual Chassis members. PR1255542
- Guest network functions (GNFs) in a node-slicing setup currently do not support Junos snapshot/recovery mechanisms. PR1268943
- The performance of an X710 NIC is lower compared to that of an 82599 NIC. A 40G line rate can be achieved at 512-byte packet size for the X710 NIC compared to 256 bytes for the 82599 NIC. PR1281366
- If a VM host snapshot is taken on an alternate disk and there is no further VM host software image upgrade, the expectation is that if the current VM host image gets corrupted, the system will boot with the alternate disk so that user can recover the primary disk to restore the state. However, if the host root file system is corrupted, the node boots with the previous VM host software instead of booting from the alternate disk. PR1281554
- Due to a vendor code limitation, ungraceful removing of summit MACsec TIC from the chassis might cause a crash or an unpredictable result. PR1284040
- Junos OS releases with a fix committed in Junos OS Releases 15.1R5-S4, 16.1R4-S3, 16.1R5, and 17.3R1 with XM-based linecards (MPC3E/4E/5E/6E/2E-NG/3E-NG) might report **DDR3 TEMP ALARM** chassisd error log message. PR1293543
- The Routing Engine gets stuck and boots from the other SSD after vmhost reboot. You must boot the Routing Engine from the primary SSD. PR1295219
- The **show dynamic-tunnels database summary** command might not show an accurate tunnels summary during the time the anchor Packet Forwarding Engine linecard is not in up state. As a workaround, use the following commands: **show dynamic-tunnels database** and **show dynamic-tunnels database terse**. PR1314763

- The chained-composite-nexthop configuration statement does not bring in a lot of gain, because TCNH is based on ingress rewrite premise. It works fine when the configuration statement is not used. PR1318984
- In JDM (running on the secondary server), the jdmd daemon might generate core files if the GNF add-image is aborted by pressing CTRL+C. PR1321803
- BGP signals tunnels are always next-hop-based tunnels. The GRE tunnels created dynamically by a BGP signal are always next-hop-based tunnels, even if the user has configured the static tunnels created by GRE to use the logical interface base. PR1322941
- When FPC restarts with Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections. If the integrity of the TCP connection fails, a TCP connection timeout occurs because of jlock hog crossing boundary value (5 seconds), causing bad consequences in the MX Series Virtual Chassis. PR1332765
- The output of the CLI command **show class-of-service fabric statistics** now includes traffic that was dropped because of internal errors in the drop counts. PR1338647
- In an EVPN-VXLAN scenario, when moving hosts between two multi-homed (MH) interfaces. During IP/MAC movement, old MAC+IP entries were deleted from the global DB, but retained in one of the local DBs. This might cause l2ald core. The normal condition is that both entries should have been deleted. PR1339543
- MACsec sessions might not get established when FPCs continuously go offline or online more than 10 times followed by restarting dot1xd. PR1344358
- The first packet pertaining to the J-Flow Packet Forwarding Engine sensor in UDP mode is missing after a line card reboot on MX150 platform. PR1344755
- During a unified ISSU that warrants a host upgrade, if the router is configured with 8 million IPv4 or IPv6 routes or more, the unified ISSU might fail, resulting in an FPC restart. PR1348825
- In some cases, online insertion and removal (OIR) of a MIC on an FPC can lead the traffic destined to the FPC to be discarded without notification. The only way to recover from this is to restart the FPC. The issue will not be seen if you use the corresponding CLI commands to take the MIC offline and then bring it online. PR1350103
- On all Junos platforms, licenses might not take effect after successfully committing a license key configuration. PR1350302
- Interface range is not supported for channelized interfaces on the EX9253. The user has to configure the interfaces individually. PR1350635
- During stress conditions, error log messages regarding route add, change, or delete might be incorrect. PR1350713
- When performing unified ISSU to the 18.2+ builds with 1334612 fix, both the From and To builds should have 1334612 fix. Otherwise CRC errors will be seen. PR1353911

- The Craftd messages are generated on MX10003 and MX204 platforms. MX10003 and MX204 routers do not have a Craft interface. Hence, these errors are expected, and can safely be ignored. When the craftd daemon tries to open the device, it fails with a junk char in the fatal error message because the error number is not mapped to a string in the kernel code. The following error messages are displayed:**Feb 20 01:49:38 MX craftd[xxxx]: craftd detected platform mx10002 Feb 20 01:49:38 MX craftd[xxxx]: LIBJSNMP_SA_IPC_REG_ROWS: ns_subagent_register_mibs: registering 1 rows Feb 20 01:49:38 MX craftd[xxxx]: fatal error, failed to open smb device: ,JÎÈ**.PR1359929
- Some of the exported packets for the sessions sensor could get fragmented. Due to this, at times the collector receives only the telemetry header part and not the payload. PR1364288
- Syslog is updated when the user tries to configure an extended packet numbering (XPN) cipher over a non-xpn supported platform such as MIC-MACSEC-20G even though commit goes through. PR1367722
- When you swap a Virtual Chassis of QFX5100 to the EX9253 for testing some heavy multicast, even when the IRB interface comes up, traffic drop might be observed. PR1369099
- When the FPC is booting up (either during unified ISSU, router reboot, FPC restart), i2c timeout errors for the SFP can be noticed. These errors are seen as i2c action is not completed because device was busy. Once the card is up, all the i2c transactions to the device work correctly, so no periodic failure is observed. There is no functional impact and these errors can be ignored. PR1369382
- The voltage high alarm might not be cleared when the voltage level comes back to normal for the MIC on MPC5E. PR1370337
- Every L2BSA subscriber creates 2 interfaces, DVLAN and RTSOCK with the same subunit (same interface name). Initially, the CLI output for **show interfaces extensive** displayed the filter information on both the DVLAN and RTSOCK interfaces. Functionally, the filter information should only be displayed on the DVLAN interface. PR1372527
- When the MIC-MACSEC-20G is in offline state after FAKE-KATS initiation, the MIC has to be brought up by issuing **chassisd** restart. Attempting to bring the MIC online via the CLI could cause the MIC to go to a hardware error state. PR1374532
- Log messages continuously appearing in the MPC console indicate the presence of a faulty SFP/SFP+, which is causing I2C transaction from the main board CPU. There is no software recovery available to recover from this situation. These logs also indicate potential I2C transaction failure with any of the 10 ports available with GMIC2 in PIC 0 resulting in unexpected bbehaviors such as , such as the link not coming up or the MIC itself not booting up on restart. **I2C Failed device: group 0xa0 address 0x70Failed to enable PCA9548(0x70):grp(0xa0)->channel(0)mic_sfp_select_link:MIC(0/0) - Failed to enable PCA9548 channel, PCA9548 unit:0, channel ID: 0, SFP link: 0mic_sfp_id_read: Failed to select link 0** To recover from these failures, is to detect and replace the faulty SFP/SFP+ plugged into the GMIC2 ports. PR1375674
- When the pfe_disable action is triggered, all the physical and logical interfaces for that Packet Forwarding Engine should be disabled. However, only physical interfaces go down, leaving the logical interfaces still in 'UP' state. PR1380784

- In rare situations at heavy traffic loads, the input frame check sequence counter might get incremented. PR1383009
- Users can still issue the command **set vmhost...** although **"permissions system-control** is not configured on system class. PR1383706
- Commit should not be allowed if you are trying to delete the **physical-cores** knob. However, there is no functional impact. PR1384014
- You can configure the purge timeout of programmable RPD clients to **never**. This means that the routes added by PRPD clients will not be deleted when the client disconnects. They will stay until the routing daemon restarts or it is deleted by the client that added the route. This can be configured by using the following CLI command: **set routing-options programmable-rpd purge-timeout never**. Note: The programmable API for setting purge timeout does not support this feature yet. PR1384303
- If using VRF configurations along with a static default route to em0.0, the interface flaps might result in RPD end up. PR1386475
- In low-end 32-bit systems, rpd has a lower level of available memory. It is desired to have a log message to alert cusers when the average memory usage or transient memory usage exceeds thresholds. PR1387465
- During the Zero Touch Provisioning (ZTP) process, the default route is cleaned up by code. As a result, if a static default route is configured in the initial configuration (configuration file downloaded from the file server for ZTP), the route will fail to work. This might lead to ZTP failure or a device access issue after ZTP. PR1387724
- On an MX platform enabled with enhanced subscriber management, if the filter service is enabled for each subscriber, and a large scale of broadband edge (BBE) subscribers (for example, 10000) are logging in and out repeatedly, the FPC might crash due to this rare issue. PR1388120
- In cases of PS over RLT at high scale, removing and adding back CoS configuration can cause the FPC to enter a hard error state. PR1388487
- In a Junos Fusion Provider Edge (MX Series) scenario, all the FPCs might restart after committing the changes to the VLAN/encapsulation on the extended port if the parameter **per-interface-per-member-link ingress** is configured for sourced routing statistic by using the command **set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress**. PR1392071
- An MPC card with AFEB or TFEB with channelized OC MIC might crash and generate core files. PR1396538
- The Junos rpd daemon has facilities to attempt to trap certain classes of nonfatal bugs by continuing to run, but leaves a "soft" core file. Leaving a soft core file is intended to be nondisruptive to routing and forwarding. Users can disable these files from being generated. PR1396935
- The router advertises the ESMC quality level of primary reference clock even though the current clock status is holdover. This issue is applicable to all platforms. PR1398129

- When the MoFRR feature is used in a scaled environment (in terms of number of routes and next hops), the actual convergence of multicast traffic might reach hundreds of milliseconds because of sub-optimal handling of MoFRR forwarding states on the Packet Forwarding Engine level. PR1399457
- The **ether-pseudowire zero-control-word** configuration option under the **forwarding-options enhanced-hash-key family mpls** statement does not take effect in a Junos Node Slicing setup. Although configured as: **set forwarding-options enhanced-hash-key family mpls ether-pseudowire zero-control-word**. The parameter is not passed to MPC9E line card. This can impact load balancing over Abstract Fabric (AF) interface when using Pseudowire Headend Termination (PWHT) in a guest network function (GNF). PR1400881
- In a BGP-PIC instance, if a router R1 resolves on top of a multipath-route R2, where R2 has primary and backup indirect next hops, the results will be better if the backup leg is not used for resolution of R1. There is no impact on any existing CLI commands. The backup path is never used when the primary path is available. PR1401322
- The authentication module for JET RPCs and Telemetry fails in authenticating usernames or passwords of certain lengths. Hence, users will be unable to execute JET APIs or Junos Streaming Telemetry. PR1401854
- After upgrading Junos to Junos OS Release 17.2 or later releases, the statement **chained-composite-next-hop ingress l3vpn extended-space** cannot be configured any more on a logical system. PR1402390
- With the initiation of image installation on the base System of a setup with node slicing enabled, session gets terminated unexpectedly. PR1402643
- The issue happens because of the usage of incorrect free of community reference count by policy module. The issue happens in the condition wherein the BGP import policy has community actions and forwarding-table also has a policy with community actions (such as 'community add') when routes are being added. PR1406357
- The process rpd might crash after a non-forwarding route (for example, a route to an indirect next-hop association is a non-forwarding indirect next hop) that is received from multiple protocols is resolved again by using the non-forwarding path. PR1407408
- If generic routing encapsulation (GRE) over GRE tunnel is used for sending Routing Engine originating traffic, the traffic cannot be encapsulated properly although the GRE over GRE tunnel works for transit traffic. PR1411874
- A small number of tunneled subscribers might be terminated during unified ISSU to Junos OS Release 19.1R1 software due to momentary loss of IP connectivity between the LAC and LNS devices. PR1412818
- In the subscriber environment, if the client profile has no filters while the service profile has filters, after a subscriber login, the ifstate compression might be seen when deleting the current filters and then adding a different filter. When this occurs, the firewall filter might be corrupted. PR1414706
- A small number of tunneled subscribers might be terminated during unified ISSU to Junos OS Release 19.1R1 software due to momentary loss of IP connectivity between the LAC and LNS devices. PR1414928

- PCE-initiated LSPs are deleted from PCC if the PCEP session goes down and gets reestablished within the **delegation-cleanup-timeout** period. PR1415224
- A user can configure a template in the router and map that template with an external controller. The router inherits the required configuration from the template and then provisions the external controller initiated LSP. Unbinding the template from the external controller or changing the template configuration might delete the PCE initiated LSPs (only LSPs which are using that particular template). Later, the LSPs are reprovisioned by the external controller. PR1421093
- MX204 supports SFP "SFP-1GE-FE-E-T" from some releases. I2C read errors are seen when an SFP-T is inserted into a disabled state port, configured with the **set interface <*> disable** CLI command. **[M LOG: Err] smic_mx1ru_8xsfpp_mpcs_i2c_read: - SFPP set start_addr failed [M LOG: Err] I2C Failed device: group 0x812 address 0x56 [M LOG: Err] mpcs_i2c_single_io: MPCS(0) ctlr 2 group 2 addr 0x56 prio 1 flags 0x0 failed status 0x1 [M LOG: Err] smic_mx1ru_8xsfpp_mpcs_i2c_read: - SFPP set start_addr failed [M LOG: Err] I2C Failed device: group 0x812 address 0x56 [M LOG: Err] smic_sfpp_ext_phy_get_linkstate: SMIC(0/1) - SFPP ext phy read failed [M LOG: Err] smic_phy_periodic DFE tuning failed for xe-0/1/2 [M LOG: Err] smic_periodic_raw: SMIC(0/1) - Error in PHY periodic function.** PR1423858
- On vBNG platforms, if hierarchical-scheduler maximum-hierarchy-levels 3 is configured in the interface statement, this feature works when schedulers are configured on the subscriber and interface-set/ifd levels. If subscriber ifl scheduler is not configured, and only ifd scheduler is configured, traffic among subscribers will work as expected. As a workaround, configure hierarchical-scheduler maximum-hierarchy-levels 2 if subscriber ifl scheduler is not required/configured. PR1425755
- RPSD might generate a core file when **rpsd/host-route-generration** is unconfigured at the same time that reconfig toggles **set system processes routing force-64-bit**. When the issue is hit, the result temporarily impacts rpsd only. RPD is not affected due to RPSD's long time to terminate. The functional impact is that rpsd cannot be reconfigured until the previous instance finally forcibly terminates after 10 minutes, which will generate an rpsd core file. PR1429770

Infrastructure

- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. PR1315605
- This is a BIOS firmware issue and does not seem to impact any functionality. All systems running BIOS version earlier than 1.1 are reporting a warning message. As a workaround, upgrade the BIOS firmware on the devices. You can check for the firmware version on the device by querying the sysctl hw.re.biosversion. The version should be later than 1.1 for this warning to be resolved. PR1345166
- Junos OS can stop responding trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. PR1359339

Interfaces and Chassis

- In Junos OS BNG solutions, after a commit event, if the configuration contains duplicate VLAN IDs configured on aggregate and demux interfaces, the MX Series router might go into data base prompt mode and the kernel might generate core files. PR1274038
- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after upgrade. This is because of the old version of **/var/db/cfm.db**. PR1281073
- The aggregated Ethernet speed calculation changes according to 10 Gigabit Ethernet after GRES. PR1326316
- There might be a memory leak on tarnsportd when bulk SNMP polling is done on large-scale Iinterfaces and a large number of traps is created due to interface flapping. The memory leak could cause the transportd to consume high CPU for a prolonged period. PR1398967
- Static demux0 logical interfaces do not come up after a configuration change if the underlying interface is et (100-Gigabit Ethernet). After the configuration change the et interface gets flushed in order to reparse the configuration. During this, the dcd fails to create the dependency between demux0 logical interfaces and the underlying et interface, which results in flushing of the demux0 logical interfaces. This issue is seen only if the underlying interface is et. For all other interfaces this issue has been already addressed. This is a day one issue. As a workaround, either restart dcd or (reboot the Routing Engine), to clear the problem or else use **commit full** instead of **commit** while committing a new configuration. PR1401026
- On MX Series platforms, EX-SFP-1FE-LX SFP does not initialize with MIC-3D-20GE-SFP-E(EH). PR1405271
- When an unnumbered interface is binding to an interface that has more than one IP address and one of the IP addresses is deleted, the family inet of the unnumbered interface might be deleted. The issue results in traffic loss for all the services that rely on the family inet of the unnumbered interface. Configure **preferred-source-address** on the unnumbered interface to prevent deletion of the IP address, thereby avoiding the deletion of the family inet of the unnumbered interface. PR1412534

Layer 2 Features

- This issue occurs in routers equipped with the line cards T4000-FPC5-3D, MX-MPC3E-3D, MPC5E-40G10G, MPC5EQ-40G10G, and MPC6E MX2K-MPC6E, If the router is working as a VPLS PE device because of MAC aging every 5 minutes, a VPLS unicast traffic is flooded as unknown unicast every 5 minutes. PR1148971
- On all Junos with Trio platforms, the unicast traffic might get dropped when it is passed from an Integrated Routing and Bridging (IRB) interface towards label switch interface (LSI) if the aggregated Ethernet (AE) load balancing adaptive or per-packet is configured. PR1381580

- In an LDP-VPLS setup where user-defined mesh groups are configured in a VPLS instance and the LDP-VPLS also has at least one directly connected CE interface configured under the instance, if all directly connected CE interfaces go down, the pseudowire for that instance will be transited to ST state and RS state. This might cause the traffic loss for one CE site to peer CE site. If the knob **connectivity-type permanent** is configured, this issue will not be observed as the instance will remain in UP state. PR1415522

MPLS

- When using **mpls traffic-engineering bgp-igp-both-ribs** with both LDP and RSVP enabled, Constrained Shortest Path First (CSPF) for interdomain RSVP LSPs cannot find the exit area border router (ABR) when there are two or more such ABRs. This causes interdomain RSVP LSPs to break. RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. PR1048560
- When **minimum-bandwidth** and **bandwidth** commands are present in the configuration, the bandwidth selection of the LSP is inconsistent. PR1142443
- In a CE-CE setup, traffic loss might be observed over the secondary LSP on primary failover. PR1240892
- If the primary link goes down immediately after bypass (for example, FPC containing both primary and bypass or, both primary & bypass FPCs go down simultaneously) such that primary link goes down even before the PLR sends out any path message after bypass is down, then the nodes downstream of the PLR along the LSP path will be left with stale LSP state until refresh timeout. This condition will not result in any traffic loss. PR1242558
- With nonstop active routing (NSR), when the routing protocol process (rpd) restarts on the master Routing Engine, the rpd on the backup Routing Engine might restart. PR1282369
- In case of CSPF-disabled LSPs, if the primary path ERO is changed to an unreachable strict hop, sometimes the primary path stays up with the old ERO. The LSP does not switch to standby secondary. PR1284138
- For an SR-TE path with "0" explicit NULL as the innermost label, the SR-TE path does not get installed with label "0". PR1287354
- For static short reach traffic engineering (SR-TE), the binding SID entry disappears after modifying binding (swapping) SID values for two SR-TE LSPs. As a workaround, delete the BSID->P1 and create BSID->P2. PR1289950
- Executing a **restart chassisd** in an MX Series Virtual Chassis router with the following elements configured might result in a core file.
  - IGP OSPF/OSPF3 (area 0, LFA) ISIS (level 2, LFA) LDP synchronization ipv4 and ipv6
  - IBGP dual, redundant route reflection IPv4 and IPv6
  - MPLS LDP (IGP synchronization, track IGP metric) RSVP (node link protection, adaptive, auto bandwidth, refresh reduction)

  - L3VPN OSPF OSPF3 BGPv4 BGPv6 RIPv2 static MBGP NGEN-MVPN l3vpn cnh with ext space any to any hub and spoke MPLS access Ethernet access multicast extranet per vpn and per prefix labels SRX based network address translation SRX based firewall
  - Direct Internet Access EBGP
  - CoS BA/MF classification policing/shaping queuing/scheduling hierarchical queuing/shaping/scheduling 8 traffic classes
  - BFD/OAM/CFM liveness detection
  - Load Balancing L2 aggregate ethernet IP equal cost multi path MPLS equal cost multi path
  - High Availability GRES/NSR ISSU fabric redundancy tail end protection BGP prefix independent convergence edge
  - Security loopback filter arp policers control plane traffic policers urpf check with all feasible paths ttl filtering jflow/ipfix export only SRX based DDOS

  PR1352227

- Traceroute MPLS from Juniper Networks to Huawei routers does not work as expected due to unsupported TLV. PR1363641
- When performing traceroute to a remote host for an MPLS LSP using the command **traceroute mpls bgp**, in very rare cases, the mplsoam daemon might hold the stale BGP instance handle in the query to the rpd process to get the information for the Forwarding Equivalence Class (FEC). As a result, rpd crash might occur because of the invalid instance and cause traffic impact till rpd comes back up. PR1399484
- When a make-before-break (MBB) new instance signaling experiences an error and before retry is finished, other triggers such as auto bandwidth adjustment timer expiration have to be blocked until MBB finishes. Once the MBB finishes instance switching, a blocked trigger needs to be scheduled, but it should only be triggered after optimize-adaptive-teardown timer expires. In the affected releases, the blocked trigger is scheduled immediately after instance switching without taking the optimize-adaptive-teardown timer into account. This causes old instance to be torn down before the whole system finishes changing routes using the new instance, which leads to traffic loss. PR1402382
- On Junos OS platforms with scaled MPLS labels used, when the system is already running with high load, inefficient labels allocation might cause even higher CPU utilization at 100% for hours. The issue might affect traffic. PR1405033
- Dynamically configured RSVP LSPs for LDP link protection might not come up after disabling/enabling protocol mpls. PR1432138

Network Management and Monitoring

- The snmpd daemon leaks memory in snmpv3 query path and crashes. The issue is caused by a memory leak when the request PDU is dropped by SNMP when the configuration **snmp filter-duplicates** is enabled. Each request PDU has a structure pointer for the SNMPv3 security details. This is allocated when the PDU is created or cloned. But while dropping the duplicate requests the structure is not freed, which causes the memory leak. PR1392616

Platform and Infrastructure

- The issue occurs when 120 bridge domains (among a total of 1000 bridge domains) have XE/GE links toward the downstream switch and LAG bundles as uplinks toward upstream routers. The XE/GE link is part of the physical loop in the topology. Spanning-tree protocols such as VSTP, RSTP, and MSTP are used for loop avoidance. Some MAC addresses are not learned on DUT when LAG bundles that are part of such bridge domains are flapped along with other events such as a spanning-tree root bridge change. PR1275544
- An accuracy issue occurs with three-color policers of both type single rate and two rate in which the policer rate and burst-size combination of the policer accuracy vary. This issue is present starting in Junos OS Release 11.4 on all platforms that use MX Series ASIC. PR1307882
- When chassis control restart is done with the CoS rewrite rule configured on the aggregate Ethernet interface, **Platform failed to bind rewrite** messages might be seen in syslog. The issue is specific to aggregate Ethernet interfaces. It is a timing issue that can occur when logical interfaces deletion is delayed due to high scale. When logical interfaces come up again after restart, they have different indices. PR1315437
- Provides ability to configure host rsyslog from a Junos guest.

  HOST side: The facility is one of the following keywords: auth, authpriv, cron,daemon, kern, lpr, mail, mark, news, security (same as auth), syslog,user, uucp and local0 through local7. The keyword security should not be used anymore and mark is only for internal use and therefore should not be used in applications. Anyway, you may want to specify and redirect these messages here. The facility specifies the subsystem that produced the message, that is all mail programs log with the mail facility (LOG_MAIL) if they log using syslog. The priority is one of the following keywords, in ascending order: debug, info, notice, warning, warn (same as warning), err, error(same as err), crit, alert, emerg, panic (same as emerg). The keywords error, warn and panic are deprecated and should not be used anymore. The priority defines the severity of the message.

  Guest side: **https://www.juniper.net/documentation/en_US/junos/topics/reference/general/syslog-facilities-severity-levels.html**

  remote : sync syslog server configuration from Junos to Linux & modify rsyslog.conf

  set vmhost/app-engine syslog host any any

  set vmhost/app-engine syslog host match xxx. PR1341549

- This is a minor enhancement to add a UI to copy files from a Junos VM to a Linux host. PR1341550
- In a filter list (input-list/output-list) scenario, when the filters in the same filter list refer to a same nested filter, the FPC might crash continuously. The issue results in traffic loss during FPC crash and reboot. PR1357531
- In a Layer 3 VPN topology, traceroute to a remote Provider Edge (PE) device for a CE-facing network results in an ICMP TTL expired reply with a source address of only one of the many Customer Edge-facing networks. In Junos OS Releases 15.1R5, 16.1R3, and 16.2R1 and later releases, there is a kernel sysctl value, icmp.traceroute_l3vpn. Setting this to 1 will change the behavior to select an address based on the destination specified in the traceroute command. This PR adds the option to the configuration. PR1358376
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur. Even with enhanced convergence configured, there is no guarantee that subsecond convergence will be achieved. There will be no guarantee that sub-second convergence will be achieved. PR1371493
- Validation for RFC 2544 feature (test start conditions) on MX Series routers is broken in Junos OS Releases 18.3R1 and 18.4R1. An invalid test start condition could lead to an inconsistent state between the Routing Engine and the Packet Forwarding Engine. PR1396751
- FPC reset might be observed in the following scenarios on a scaled setup –IGP flap (carrying multiple LSPs) and/or clearing multiple MPLS LSPs or any similar event causing churn in router. PR1398502
- In some cases, PS interfaces over RLT might be shown as 'up' but might not pass traffic. Log messages reporting ASIC errors and a chassis alarm reporting hard FPC errors might also be seen. PR1400269
- On MX Series platforms in a VPLS scenario, when the **interface-mac-limit packet-action-drop** knob is configured, in case of MAC moves, the new MAC might not be learned because of a race condition resulting from an unusual update of the learn limit in the Packet Forwarding Engine (the HW learn limit counter displays unexpected behavior and increases to a huge and negative number). This can result in a packet drop. PR1410162

Routing Policy and Firewall Filters

- The rpd might crash during the policy configuration changes. PR1357802

Routing Protocols

- When only the default routing instance is present, the Junos **show bgp summary** command does not show the BGP ESTABLISH state. If the BGP state is not an ESTABLISHED state, then it shows the states as design (that is, Active, Idle, or Connect). If there is a routing-instance configured (apart from master routing instance inet.0), the BGP ESTABLISH state is shown correctly. The issue occurs for IPv4 BGP sessions only; on IPv6 all the BGP states as always shown as default. PR600308
- In rare cases, the rpd might generate a core file with the error **rt_notbest_sanity: Path selection failure on ...**. The core file is soft, which means there should be no impact to traffic or routing protocols. PR946415
- JTASK_SCHED_SLIP for rpd might be seen when routing is restarted or when the OSPF protocol is disabled with scaled BGP routes in an MX104 router. with scaled BGP routes in an MX104 router. PR1203979
- LDP OSPF are 'in sync' state because of IGP interface down with LDP synchronization enabled for OSPF. **user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.01Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050Adj count: 1Hello: 10, Dead: 40, ReXmit: 2, Not StubAuth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTCProtection type: NoneTopology default (ID 0) -> Cost: 1050LDP sync state: in sync, for: 00:04:03, reason: IGP interface downconfig holdtime: infinity**. 'IGP interface down' is observed as the reason because although LDP notified OSPF that LDP synchronization was achieved, OSPF did not take note of the LDP synchronization notification, because the OSPF neighbor was not up yet. The issue is under investigation. PR1256434
- There are scenarios in which the application allocates and caches next-hop templates. This causes next-hop template cache to grow continuously. But when the application clears the local cache, the memory is freed to next-hop template cache. But the next-hop template cache does not have the code to shrink the cache and free memory back. So the next-hop template memory is trapped in the cache and cannot be used for other purposes. But if the same BGP routes and next hops come up again, they will reuse the templates from the cache and not consume additional memory. PR1346984
- A guest network function (GNF) with rosen6 [rosen 6?] multicast might display stuck KRT queue entries after recovery from a dual Routing Engine reboot at the BSYS. PR1367849
- When the loopback interface is configured in a logical-system and Routing Engine-based micro BFD is configured to use the loopback address as source address, BFD packets go out with source address belonging to outgoing interface rather than the loopback address. Due to this issue, the micro BFD session might not be able to come up. PR1370463

- On all Junos OS platforms with BFD for the static route configured, when the BFD session is brought down by changing the VLAN ID of the local interfaces, the static route might persist in the routing table. PR1385380
- An invalid label operation might occur for OSPF segment routing one-hop neighbors connected through an unnumbered interface. PR1386133
- At scale, a guest network function (GNF) with PS over RLT and multiple MPCs might show BFD flap at recovery. PR1386574
- In a BGP scenario with multipath enabled, if you apply import or export policy of IPv6 routes with an IPv4 next hop to a BGP neighbor, the rpd might crash continuously. PR1390428
- If an import policy is applied to a BGP neighbor and the policy has an indirect IPv4 next hop for IPv4 and IPv6 routes (IPv6 routes resolved over IPv4), when the unresolved BGP route is withdrawn, rpd crash might be seen. PR1391568
- The route selection mechanism of rpd has multiple user-configurable mechanisms by which route ordering might be changed. To assist with debugging issues with defects in the route selection code, rt_notbest_sanity() was created as a function that would generate a low-priority soft core that did not let rpd to crash when route selection was incorrect. However, there have been circumstances wherein not-best was incorrectly being determined. This PR addresses one such situation where routes are learned or redistributed from non-BGP protocols and had an AS_PATH attribute. Using BGP route selection rules, if a BGP route and a non-BGP route had a leading AS_PATH with the same AS, BGP MED selection rules for grouping were being applied. Such MED selection should only be done using BGP-only routes. Such a situation can arise from various BGP-carried VPN protocols wherein routes from the VPN protocol generated IPv4 routes are redistributed from one routing instance to another. An example of this would be an EVPN route. PR1391767
- The **as-path-group** configuration is limited in scale. With 10,000 lines, scheduler slips are seen, impacting other work rpd is doing such as protocol keep-alives. To avoid the scheduler slips (CPU exhaustion), change how the as-path-group is structured. The issue occurs due to two factors: the number of **as-path** statements under the **as-path-group** and the wildcards in each of these. In this PR, there is a new Junos knob introduced: **set policy-options asregex-optimize**. The default feature is **no-optimize**. PR1396344
- This issue occurs on Junos OS platforms that have Multicast Only Fast Reroute (MoFRR) and Join Load Balance (JLB) automatic features enabled and configured in a scaled setup. if it is configured in scaled setup. When the active reverse path forwarding path is disabled by some operations, for example, the metric of the active interface is increased to make it not be active anymore, there might be unexpected packet drop for about 5 seconds due to this timing issue. PR1401802
- In a multicast routing scenario using PIM, if the static route is configured with qualified-next-hop for the multicast source, the rpd process might crash. This is because the qualified-next-hop points to the GF_DLI (Gateway Family Data Links) address, which PIM is unable to process, resulting in the crash. PR1408443

- In a BGP scenario with an indirect next hop, if uRPF or route record is enabled, and then BGP multipath is enabled, a background job loop might be formed and the CPU utilization of the rpd process might be stuck at 100%. PR1414021
- In Layer3 VPN scenario with multipath enabled for BGP Layer3 VPN family, if the knob **no-vrf-propogate-ttl** and **maximum-prefix** are configured for VRF, in some certain conditions, the rpd might crash when the maximum-prefix is hit and the withdrawal of VPN route occurs. PR1427147
- In IS-IS IPv6 scenario, if MTU is changed under the Logical Interface (IFL) level for family inet6, the IS-IS IPv6 route might be deleted and might not be reinstalled. These routes remain present in IS-IS database and IS-IS adjacency remains UP as well. The reason is that IS-IS interface data is not added for IPv6 unicast topology after the interface MTU changing event and this does not allow the IS-IS IPv6 routes to get resolved. PR1420776

Subscriber Access Management

- Sometimes, when PPPoE subscribers log in and log out from Junos OS Release 16.1 releases, the following messages are generated: **user@devcie> show log messages | match authd authd[5208]: sdb_app_access_line_entry_read_by_uifl: uifl key 'demux0.xxxxxxxx': snapshot failed (-7) authd[5208]: sdb_app_access_line_entry_read: uifl key 'demux0.xxxxxxxx': read failed**. These messages indicate that the authd daemon for subscriber authentication is attempting to read private data for an underlying interface that no longer exists (-7 = SDB_DATA_NOT_FOUND). These messages have no impact and can be safely ignored, where the authd daemon is asking sdb for a record that no longer exists. PR1236211
- authd reuses addresses too quickly before jdhcpd has completely cleaned up the old subscribers, which leads to error log flooding. The log shows: **jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815**. PR1402653

User Interface and Configuration

- The test configuration **/config/rescue.conf.gz** fails commit check for the dynamic profile when the subscriber is active. PR1376689

VPNs

- The multicast VPN MIB was not being properly compiled into the Juniper MIB package bundle. Mib-jnx-mvpn.txt needs to be included as part of the Juniper Enterprise MIB set. PR1394946
- In a segmented inter-AS NG-MVPN scenario, when the PE router receives a C-multicast (or leaf AD) route with more than one community from a remote AS, the route might be rejected due to incorrect route-target community matching. PR1405182
- In a rare scenario with a multicast extranet VPN, rpd can crash because the reference count of the next hop becomes 0. PR1419891

SEE ALSO

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in the Junos OS 18.3R2 Release for MX Series routers.

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

Resolved Issues: 18.3R2

Application Layer Gateways (ALGs)

- DNS requests with EDNS options might be dropped by DNS ALG. PR1379433

Authentication and Access Control

- MAC move might occur in a DHCP security scenario. PR1369785
- The dot1xd might crash when dot1xd receives incorrect reply length from the authd. PR1372421
- Push-to-JIMS now supports pushing the authenticated entry to all online JIMS servers. PR1407371

Class of Service (CoS)

- FPC card might reboot when changing CoS mode from hierarchical-scheduler to per-unit-scheduler. PR1387987

- The cosd process might crash after commiting configuration changes through netconf. PR1403147

EVPN

- EVPN type-5 route might be lost if **chained-composite-next-hop** is configured. PR1362222
- Packet drop is seen in EVPN stitching with IRB configured. PR1363935
- The EVPN implementation does not follow RFC-7432. PR1367766
- Small rpd memory leak is seen when configuring EVPN. PR1369705
- EVPN A/A multihomed PE device occasionally prefers to route to a directly connected prefix using LSPs toward the multihomed peer instead of going directly out the IRB interface (which is up). PR1376784
- In an EVPN A/A scenario with an MX Series router or an EX Series switch acting as the PE device, flood next hops to handle BUM traffic might not get created or miss certain branches when the configuration is performed in a particular sequence. PR1377749
- The RA packets might be sent out without using the configured virtual gateway address. PR1384574
- A few minutes traffic loss might be observed during recovery from link failure. PR1396597
- The BUM traffic might not be flooded in an EVPN-MPLS scenario. PR1397325
- The IPv6 link-local address for virtual-gateway address is marked as duplicate in EVPN. PR1397925
- When committing a configuration while adding a VLAN to an EVPN instance and an aggregated Ethernet interface respectively the newly added VLAN interface count might be zero (0) in that bridge domain. PR1399371
- EVPN type 2 MAC+IP route is stuck when the route advertisement has 2 MPLS labels and withdrawal has one label. PR1399726
- The rpd core file is generated upon Routing Engine switchover with scaled EVPN configuration. PR1401669
- The rpd crashes due to memory corruption in EVPN. PR1404351
- EVPN database and bridge MAC table are out of sync post core link flap. PR1404857
- The rpd might crash on a leaf node when handling the withdrawal of remote or local MAC address in an EVPN-VXLAN scenario. PR1405681
- The rpd might crash after NSR switchover in an EVPN scenario PR1408749
- The next hop is not cleaned up properly when one of the multi-homed CE-PE links goes down. PR1412051

Forwarding and Sampling

- LTS subscriber statistics is reporting to RADIUS. PR1383354
- Adjusting the **mac-table-size** configuration might cause l2ald crash. PR1383665
- The LSI binding for the IPv6 neighbor is missing. PR1388454

- The filter counter is not written to the accounting file when accounting is enabled on the bridge firewall filter. PR1392550
- The l2ald process might crash when doing **commit check** for some specific configurations. PR1395368
- In Junos OS Release 13.3R9.13, the firewall filter action decapsulates GRE, IP-over-IP, and IPv6-over-IP. However, in Junos OS Release 17.3R3.9, it only decapsulates GRE. PR1398888

General Routing

- Routing Engine-Packet Forwarding Engine out-of-sync errors might be seen in syslog. PR1232178
- An mspmand core file might be generated in rare conditions due to a high rate of TCP traffic. PR1253862
- Error messages might be seen if the aggregated Ethernet interface host on the MPC-3D-16XGE line card flaps. PR1279607
- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/mspsmd. PR1284654
- The **RE does not have MAC map for mac type 7** error message might be seen on MX10003 routers. PR1345637
- Large-scale users log in and log out might cause a mgd memory leak. PR1352504
- Traffic loss might be seen on the new master Routing Engine after the interface flaps followed by Routing Engine switchover in a VRRP scenario. PR1353583
- On MX Series routers, network slicing GNF is allowed to install incompatible images without warnings. PR1353773
- The packets might be dropped when they go through the MX104 built-in interface. PR1356657
- MPC/FPC might be unable to reply to request messages to the Routing Engine in a highly scaled subscriber scenario. PR1358405
- The **show chassis ethernet-switch** command output on MX-TVP platforms is different from that of the MX2010 router. PR1358853
- FPC core file might be observed after GRES switchover. PR1361015
- The MX Series router functioning as a BNG does not generate ESMC/SSM quality level failed SNMP trap. PR1361430
- On the MX10003, the alarm LED reflects stale entry on the backup Routing Engine, post GRES switchover. PR1361728
- The MS-MPC might reset continuously on MX Series routers. PR1362271
- The inline J-Flow sampling configuration might cause FPC crash on MX Series routers. PR1362887
- MX-Virtual Chassid: request to record VCCP heartbeat state change in syslog by default. PR1363565
- FPM board status is missing in the SNMP MIB walk result. PR1364246
- Netproxy service client component fails to start after issuing the **request vmhost reboot** command. PR1365664

- The following syslog errors are seen on MX960 routers: **LOG : Err] Failed to allocate 2 jnh-dwords for encap-ptr(ether-da)!,LOG: Err] gen_encap_common: jnh-alloc failed! 8** PR1366811
- When you configure VRRP delegate-processing with Apache Tomcat enabled, the Packet Forwarding Enging drops the VRRP packets and counts software error. PR1369503
- The MPC5E, MPC2E-NG, or MPC3E-NG might crash and restart during unified ISSU. PR1369635
- SNMP MIB walk causes KMD errors. PR1369938
- The rpd might crash after Routing Engine switchover is performed or the rpd is restarted if interface-based dynamic GRE tunnel is configured. PR1370174
- SFP-1FE-FX optics is not coming up on GMIC. PR1370962
- The bbe-smgd might crash when the FPC is restarted. PR1371926
- Image installation on SD fails with the **Unable to read reply from software add command to re1; error 1** error. PR1372877
- A core file is generated in ifinfo at pif_af_fe_info pif_af_ifd when displaying af interface information. PR1373436
- SFP-100BASE-BX10-U and SFP-100BASE-BX10-D are not supported on 20x1-Gigabit Ethernet and 2x10-Gigabit Ethernet MACsec MIC due to a microsemi PHY limitation. PR1373795
- LDP convergence delay might be seen after IGP metric change with **bgp-igp-both-ribs** configured. PR1373855
- Cosmetic log **warning: [---] is protected, '---' cannot be deleted** is seen after commit using **configure private** in a configuration with the **protect** flag present. PR1374244
- The filter service might fail to get installed for the subscriber in a scaled BBE scenario. PR1374248
- FPC might not be able to work properly if one child interface is removed from an aggregated Ethernet bundle in a dynamic VLAN subscriber scenario. PR1374478
- A few L2BSA subscribers might be stuck in init, terminating, or terminated status after previous log out. PR1375070
- SFB and PDM/PSU related information is missing in jnxBoxAnatomy MIB on high-end MX Series routers. (MX2010/2020). PR1375242
- The bbe-smgd core file might be seen after doing GRES. PR1376045
- MS-MPC might have performance degradation under scaled fragmented packets. PR1376060
- Interface optic output power is not zero when the port has been disabled. PR1376574
- The **Power Supply failed** trap might not be generated on MX Series routers. PR1376612
- Disabling OAM might cause the broadband edge daemon to crash. PR1377090
- Packets might be dropped on the data plane in an inline J-flow scenario. PR1377500
- MQTT keepalive timeout messages are seen in case of slow JTI collectors. PR1378587

- After NAT64 router (with MS-MPC) translates an IPv6 fragment to IPv4 fragment, the router is not inserting the right value in the identification field of the IPv4 header. PR1378818
- Traffic might get discarded without notification when CoS configuration is changed on a PS interface. PR1379530
- Protocol adjacency might flap and FPC might reboot if jlock hog occurs. PR1379657
- Remove the chassisd alarms for FPCs exceeding 90 percent of power budget and exceeding 100 percent of power budget. PR1380056
- The rpd might crash on the new master Routing Engine when performing GRES. PR1380298
- Encryption and decryption are not happening because the Packet Forwarding Engine discards it while testing that group-VPN member established using the authentication method preshared key ascii-text. PR1381316
- Traffic is discarded without notification when an FPC is taken down in an MC-LAG scenario. PR1381446
- Memory leak is observed in MS-MPC line card. PR1381469
- Constant memory leak might lead to FPC memory exhaustion. PR1381527
- Subscribers might not be able to log in after double GRES, after reboot, or after configuration. PR1382050
- On MX10003 routers running Junos OS Release 18.3R1, unified ISSU might fail if QSA is plugged in. PR1382126
- The MPC6E might crash while fetching PMC device states. PR1382182
- Flows are getting exported before the expiration of the configured active timeout value. PR1382531
- Expected **inline-ipv4-export-packet-failures** is not listed in the **show services accounting error** command. PR1382873
- MAC addresses might disappear if the interface MTU of EVPN PE device is changed. PR1382966
- The chassisd might crash due to HW-DB errors on TVP-based platforms. PR1383246
- Domain name is not reported as part of the LLDP system name in the **show lldp neighbor** command. PR1383295
- The configuration configured through NETCONF session might fail. PR1383567
- The kmd crashes with generation of core file after bringing up the IPsec connection. PR1384205
- CoS attachment might be mistakenly removed for DHCPv4 stack when DHCPv6 stack fails to be brought up for single-session dual-stack subscriber. PR1384289
- Missing **interface-description** configuration statement for static subscribers. PR1384421
- MBFD flaps because clksync congests the scheduler for 100 ms. PR1384473
- Multiple bbe-smgd core files are generated with reference to bbe_mcast_vbf_dist_policy_service_encoder( ). PR1384491
- Subscriber connection setup is 30 percent lower than expected. PR1384722

- The MPLS packets with number of labels more than 8 will not be processed by jflow. PR1385790
- On vMX, the vFPC CPU utilization is very high.PR1385853
- The device with more than 5 IP addresses configured in the DHCP server-group goes into amnesiac mode after reboot. PR1385902
- IPSec VPN traffic might fail when passing through MS-MPC of MX Series routers with CGNAT enabled. PR1386011
- Representation of memory units is changed from Gigabytes[GB] to Gibibytes[GiB] in the help string under resource template hierarchy. PR1386516
- In a subscriber management environment, DHCP subscribers might get stuck in terminated state. PR1386662
- IPv4 and IPv6 VIP routes are not withdrawn after aggregated Ethernet and VLAN with IRB flap. PR1386713
- The rpd might crash due to a memory leak issue in route resolution code paths. PR1386755
- Agent ID in the **show sflow** command is displaying lo interface IP address instead of fxp0 IP address. PR1386890
- In case an LSP is locally configured without an explicit path ERO, the object remains empty in the PCRpt generated by PCC. PR1386935
- Uninitialized EDMEM[0x400094] Read (0x6db6db6d6db6db6d) logs are seen with sampling applied to a subscriber with routing-service applied. PR1386948
- On MX2000 routers, backup CB's chassis environment status displays **Testing** after the backup CB comes online by removal or insert operation. PR1387130
- The pccd might crash when changing delegation-priority. PR1387419
- The bbe-smgd process might crash when two subscribers log in with the same framed-route prefix and preference values. PR1387690
- Output of the **show class-of-service interface** command incorrectly shows adjusting application as PPPoE IA tags for DHCP subscribers. PR1387712
- Some SFBs might go down when one of the PSMs in the chassis generates a bad output voltage which is out-of-range. PR1387737
- The bbe-smgd process generates repeated core files and stops running as a result of long-term session database shared memory corruption. PR1388867
- IPsec IKE keys are not cleared when delete or clear notification is received. PR1388290
- The bbe-smgd might not respond to the NS message for the SLAAC client on dynamic VLAN. PR1388595
- Fabric drops might be seen when using a newer generation of MPC with SFB2. PR1388780
- Incorrect value for flow packets or octets fields might be seen in an inline-jflow scenario. PR1389145

- IGMP group threshold exceed log message prints a wrong demux logical interface. PR1389457
- Excluding the **speed** CLI option under the interface level. PR1389918
- The jnxFruState might show incorrect PIC state after replacing an MPC with another MPC having less PICs. PR1390016
- CoS adjustment-control-profile configuration for application DHCP tags does not get applied. PR1390101
- Traffic destined to VRRP VIP gets dropped as filter is not updated to related logical interfaces. PR1390367
- Delay in CLI output with second or more **show subscriber <> extensive** queries occur when the first session is at the -(more)- prompt displaying the **show subscribers extensive** command output. PR1390762
- Trailing characters appear in GNMI get API reply. PR1390967
- All the BBE and ESSM subscriber sessions might be lost after GRES or unified ISSU. PR1391409
- The **routing-engine-power-off-button-disable** configuration statement does not work on MX204 and MX10003. PR1391548
- The bbe-smgd process might crash after committing configuration changes. PR1391562
- The bbe-smgd process might crash in a corner case if family inet6 is used in dynamic profile. PR1391845
- On MX2000, fans start spinning at high speed upon inserting previously offlined FPC. PR1393256
- There is a third-generation FPC reboot loop because of internal interface issues. PR1393643
- FPC might reboot on vMX in a subscriber scenario. PR1393660
- Junos OS enhancement configuration statement added to modify mcontrol watchdog timeout. PR1393716
- If FPGA on the new master CB has a specific hardware failure, the chassid might keep crashing after GRES switchover. PR1393884
- MPC7, MPC8, or MPC9 might not boot on MX Virtual Chassis. PR1396268
- The MS-MPC might generate a core file when mspmand receives a non-syn packet of TCP. PR1396785
- Enabling the Flex-Flow-Sizing takes more than 12 minutes to move to steady state. PR1397767
- The **show system errors active** command is not showing an error message for MPC3E NG HQoS. PR1398084
- Kernel core files are generated on vMX. PR1398320
- MPLSoUDP tunnels do not come up on interface route - dyn_tunnel_fwd_route_eligible because next-hop type is interface. PR1398362
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. PR1398398
- IPsec tunnel cannot be established because the tunnel SA and rule are not installed in the PIC. PR1398849
- The bbe-smgd process might crash when executing the **show pppoe lockout** command. PR1398873
- Wrong timestamp is displayed in the jvision collector log file. PR1399829

- ARP refresh functionality might fail in an EVPN scenario. PR1399873
- The mgd-API might crash due to memory leak. PR1400597
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- The authd might crash when issuing the **show network-access requests pending** command during the authd restart. PR1401249
- The **show | compare** command output on global group changes lose the diff context after a rollback or **load update** is performed. PR1401505
- The subscriber route installation fails due to some interfaces states are not properly installed. PR1401506
- FPC core files are generated due to a corner case scenario (race condition between RPF and IP flow). PR1401808
- The Framed-Route beyond the first might not be installed in a DHCP subscriber management environment. PR1401148
- Traffic loss is seen for IGMP subscribers after GRES. PR1402342
- The MPC might crash due to the CPU hogging by dfw thread. PR1402345
- Some error logs might be seen on FPC when reading is attempted from uninitialized memory location. PR1402484
- FPC might crash after you offline or online MIC-3D-16CHE1-T1-CE-H. PR1402563
- DHCP subscriber cannot reconnect over dynamic VLAN demux interfaces due to RPF check failure. PR1402674
- Host outbound traffic might be dropped on MPC7, MPC8, and MPC9. PR1402834
- Smg-service could become unresponsive when doing some GRE-related CLI operations. PR1403480
- The time synchronization through PTPoE might not work when Enhanced Subscriber Management is enabled on MX Series routers. PR1404002
- Continuous kernel crashes might be observed on the backup Routing Engine or VC-BM. PR1404038
- With MS-MPC and MS-MIC service cards, syslog messages for port block interim might show 0.0.0.0 for the private IP address and PBA release messages might show the NAT'd IP address as the private IP address. PR1404089
- The FPC might crash in a CoS scenario. PR1404325
- The repd continues to generate core files on VC-BM when there are too many IPv6 addresses on one session. PR1404358
- The **targeted-broadcast** statement does not work on IRB interfaces. PR1404442
- Configuration load override or load replace resets ANCP neighbors. PR1405318

- MPC might generate a core file after restarting FPC that belongs to targeting aggregated Ethernet and host subscribers. PR1405876
- NAT64 translation issues of **ICMPv6 Packet Too Big message** with MS-MPC/MS-PIC. PR1405882
- Fabric performance drop on MPC7, MPC8, and MPC9E and SFB2-based MX2000 routers. PR1406030
- Traffic impact might be seen if **auto-bandwidth** is configured for RSVP LSPs. PR1406822
- Layer 2 VPN will flap repeatedly after link up between PE device and CE device under "asynchronous-notification" and "some types of MICs" conditions. PR1407345
- Ephemeral database might get stuck during commit. PR1407924
- Traffic forwarding fails when crossing VCF members. PR1408058
- ToS/DSCP byte is not copied into the outer IPsec header during IP header preservation. PR1408168
- An alarm mismatch in total memory is detected after executing the **reboot vmhost both** command. PR1408480
- The MPC line cards might crash when performing unified ISSU to Junos OS Release 19.1R1 or later. PR1408558
- Python script might stop working due to **Too many open files** error. PR1408936
- On MX Series routers, service templates are not cleaned up. PR1409398
- Non-existent subscribers might appear in the **show system resource-monitor subscribers-limit chassis extensive** command output. PR1409767
- FPC might crash during next-hop change when using MPLS inline-jflow. PR1409807
- When using SFP+, the Interface optic output might be non-zero even though the interface has been disabled. PR1410465
- Traffic loss might be seen on MPC8E and MPC9E after request one of the SFB2s to go offline/online. PR1410813
- Kernel replication failure might be seen if an IPv6 route next-hop points to an ether-over-atm-llc ATM interface. PR1411376
- An rpd crash with **switchover-on-routing-crash** does not trigger a Routing Engine switchover and the rpd on the master Routing Engine goes into stop state. PR1412322
- During unified ISSU from Junos OS Release 16.1R4-S11.1 to Junos OS Release 18.2R2-S1.2, CoS GENCFG write failures are observed. PR1413297
- Broken support of [family inet6 filter] on ATM interface. PR1413663
- The user cannot enter into configure mode because the mgd is in lockf status. PR1415042
- The bbe-smgd process might have memory leak while running the **show system subscriber-management route route-type <> routing-instance <>** command. PR1415922

- The ECMP fast reroute protection feature might not work on MX5, MX10, MX40, MX80, and MX104. PR1417186
- SNMP trap message is not generated for jnxHardDiskMissing/jnxHardDiskFailed on MX10003 routers. PR1418461
- Due to a PPoE compliance issue, the MX Series router allows PPPoE session-id 65535. PR1418960
- MX Series routers might encounter CPU spikes on the service PIC when bringing up an IPsec peer against a DEP/NAT-T setup due to KMD injecting in 0.0.0.0/0 route. PR1419541
- A new tunnel could not be established after changing the NAT mapping IP address until the **IPEC SA Clear** command is run. PR1419542

Infrastructure

- The **jlaunchd: disk-monitoring is thrashing, not restarted** error might be seen. PR1380032

Interfaces and Chassis

- Momentary dip in traffic is seen when a GRES is performed. PR1336455
- The SONET interface will go down after enabling **keep-address-and-control** in a Layer 2 VPN scenario. PR1354713
- In case of MPLS, DMR packets are sent with different MPLS expiration bits if the MX Series router receives CFM DMM packets with varying expiration values on MPLS header. PR1365709
- In rare cases, L2TP subscribers might be stuck in terminated state. PR1368650
- Constant dcpfe process crash might be seen in an unsupported GRE interface configuration. PR1369757
- Unified ISSU could be aborted at **Timed out Waiting for protocol backup chassis master switch to complete** with MX Virtual Chassis confiiguration. PR1371297
- Some error logs (Tx unknown LCP packet) might be reported by bbe-smgd on MX Series routers. PR1378912
- Higher level OAM CFM between CE devices might not work in a VPLS scenario. PR1380799
- The dcd is restarted unexpectedly after committing a configuration with static demux interface stacking over ps interface. PR1382857
- The jpppd process might crash if the EPD value contains a format specifier. PR1384137
- The dcd core file can be seen after FPC restart if channelized interfaces are configured. PR1387962
- All DPCs might crash while adding or deleting a logical interface from the aggregated Ethernet bundle. PR1389206
- The interface-control process crashes and dcd does not restart after adding an invalid demux interface to the configuration. PR1389461
- Interim accounting updates might not be sent for subscribers after Junos OS selective update. PR1391011

- A dcd memory leak might be seen when committing a configuration change on the static route tag. PR1391323
- Error message might be seen if GR interface is configured. PR1393676
- The dcd crashes on deleting the subinterface from VPLS routing-instance when the same subinterface is also part of a mesh group. PR1395620
- The **MIC Error code: 0x1b0002** alarm might not be cleared for MIC on MPC6 when the voltage has returned to normal. PR1398301
- The backup Routing Engine might get stuck in amnesiac mode after reboot. PR1398445
- All dcd operations might be blocked if profile-db is corrupt. PR1399184
- Certain otn-options cause interface flapping during commit. PR1402122
- The subscriber might not be able to access the device due to the conflicted assigned address. PR1405055
- The cfmd might fail to start after it is restarted. PR1406165
- The **aaa-options** configuration statement for PPPoE subscribers does not work on the MX80 and MX104 routers. PR1410079

Layer 2 Features

- The backup VPLS router might still have MAC addresses after the primary router is rebooted and recovered in a VPLS environment. PR1356726
- The unicast traffic from IRB interface towards LSI might be dropped due to Packet Forwarding Engine mismatch at egress processing. PR1381580
- Flow label is still used by ingress PE though the Egress PE is not configured for Flow label in a VPLS multihomed scenario. PR1393447
- In a Layer 2 domain, there might be unexpected flooding of unicast traffic at every 32-40 seconds interval toward all local CE-facing interfaces. PR1406807
- When more than one site is added under **protocols vpls** in the routing instances, commit error will be seen but the commit is processed. PR1420082

Layer 2 Ethernet Services

- ZTP infrastructure scripts are not included for MX PPC routers. PR1349249
- RADIUS accounting statistics are not cleared after subscriber logout. PR1383265
- The subscriber's authentication might fail when the link-layer address encoded in the DHCPv6 DUID is different from the actual link-layer hardware address. PR1390422
- The SNMP query on LACP interface might lead to lacpd crash. PR1391545

- The **dot1xd[]: task_connect: task ESP CLIENT:...: Connection refused** log messages might be reported in Junos OS Release 17.4 or later. PR1407775
- DMAC problem of the IRB interface is seen for traffic over the Layer 2 circuit. PR1410970

MPLS

- RSVP authentication might fail between some Junos OS releases and cause traffic loss during local repair. PR1370182
- The rpd process might crash continuously if **nsr-synchronization** or **all flag** is used in the RSVP traceoptions. PR1376354
- The rpd might crash on the backup Routing Engine after switchover. PR1382249
- MPLS LSP will remains in down state due to routing loop detection after link flaps between PE router and egress PE. PR1384929
- Ingress LSPs are down due to CSPF failure. PR1385204
- Configured bandwidth 0 does not get applied on RSVP interface. PR1387277
- The bypass LSP might pass through unexpected path that includes the same SRLG as the protected down TE link. PR1387497
- The rpd process might crash repeatedly if the LSP destination address is set to be 0.0.0.0. PR1397018
- The rpd might crash when LDP route with indirect next hop is deleted. PR1398876
- A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based statistics are used. PR1401152
- High rpd usage results in routing protocols failure when doing SNMP walk of mplsXCTable. PR1402185
- Resources might be reserved for stale RSVP LSP when RSVP is disabled on the interface. PR1410972
- LDP crashes with the reason ldp_label_bind_route assert condition. PR1413231
- LDP route is not present in inet6.3 if IPv6 interface address is not configured. PR1414965
- LDP route missing in inet.3 when enabling TI-LFA node protection on LDP-SR stitching node. PR1416516

Network Management and Monitoring

- Syslog filtering (match **regular-expression** statement) doe not work if each line of **/etc/syslog.conf** is over 2048 bytes. PR1418705

Platform and Infrastructure

- MQCHIP CPQ block might report a major alarm. PR1276132
- Distributed multicast might not be forwarded to a subscriber interface. PR1277744
- The **show igmp statistics** command output does not include any statistics under interface aggregate for distributed multicast interfaces. PR1289415

- RLT subinterfaces not reporting statistics. PR1346403
- Some line cards might crash in a subscriber scenario enabled with distributed IGMP. PR1355334
- Traffic might drop on newly added interfaces on MX Series routers after unified ISSU. PR1371373
- Kernel and ksyncd core file is generated after recovering from a BSYS reboot. PR1372875
- The traffic traversing an IRB interface might not be tagged with a VLAN if the packets go through an additional routing instance. PR1377526
- FPC crash might occur after the FPC restarts. PR1380527
- IPv6 ping might fail for spine node in an EVPN scenario. PR1380590
- Packet drops on an interface if the **gigether-options loopback** statement is configured. PR1380746
- dfwd might crash with **DFWD_TRASHED_RED_ZONE** log messages. PR1380798
- Traffic loss is seen in Layer 2 VPN with GRE tunnel. PR1381740
- MAC learning might get stuck on MX Series routers with DPC and MPC. PR1383233
- Packet drops might be seen if the packet header is over 252 bytes. PR1385585
- jlock hog is reported at restart routing. PR1389809
- Individual command authorization might cause a mgd crash. PR1389944
- Traffic is dropped when passing through MS-DPC to MPC. PR1390541
- The RADIUS authentication does not work through management-instance for IPv6 family. PR1391160
- The lockout-period might not work for the user being locked out. PR1393839
- In Junos OS Release 18.4R1, after ifconfig goes down for PS logical interface, its link and admin status do not go down as expected. PR1396335
- RVT interface might start flapping. PR1399102
- On an MX204 router, when any command under the **show class-of-service fabric <>** hierarchy is executed, **COS_HALP(cos_halp_get_fabric_stats_per_pfe:3211): pfe_id 0 cchip 0** error messages are seen. PR1402377
- MAP-E for some ICMP types cannot be encapsulated or decapsulated on the SI interface. PR1404239
- Some files are missing during log archiving. PR1405903
- Abnormal queue-depth counters are seen in the **show interface queue** command output on interfaces that are associated to XM2 and 3. PR1406848
- IPv6 traffic might be dropped between a VXLAN bridge domain and IP/MPLS network. PR1407200

Routing Policy and Firewall Filters

- The **set metric multiplier offset** command might overflow or underflow. PR1349462
- The rpd process might crash if **then next-hop** is configured for LDP export policy. PR1388156

- The **as-path-expand last-as** configuration statement causes commit failure. PR1388159
- The rpd process might crash when **routing-options flow** configuration is removed. PR1409672

Routing Protocols

- BGP might not advertise routes on the existing BGP peer after adding a Layer 3 VPN instance. PR1237006
- The VRF static route might not be exported when **route-distinguisher-id** is used on RR in a BGP Layer 3 VPN scenario. PR1341720
- vFPC might continuously crash on vMX platform. PR1364624
- sBFD session flaps incrementally with 300 static SR clients configured with 100 ms as minimum-interval. PR1366124
- Ukernel memory leak and core crash is seen in a BGP environment. PR1366823
- The rpd process might crash after executing the commit configuration related to mapping-server-entry. PR1379558
- SSH is not working if **[edit system services ssh hostkey-algorithms]** is set or in FIPS mode. PR1382485
- The rpd might crash after issuing the **show route detail** operational command for RIP route. PR1386873
- Penultimate-hop router does not install BGP LU label, which causes traffic to be discarded without notification. PR1387746
- IGMPv3/MLD membership requests might not work normally. PR1389119
- Unexpected packet loss might be seen for some multicast groups during failure recovery with both MoFRR and PIM automatic MBB join load-balancing features enabled. PR1389120
- FPC might crash when BGP multipath is configured with protection. PR1389379
- BGP IPv6 routes with IPv4 next hop causes rpd crash. PR1389557
- Race condition causes all the BGP sessions to flap after an NSR switchover. PR1391084
- The ppmd on the Routing Engine might run with high CPU utilization after a Routing Engine switchover. PR1392704
- The rpd generates core file on the backup Routing Engine during neighborship flap when using authentication key with size larger than 20 characters. PR1394082
- The rpd process might crash when **rp-register-policy** is configured with more than 511 terms. PR1394259
- The best and the second-best routes might have the same weight value if BGP PIC is enabled. PR1395098
- BGP DMZ LINK BANDWIDTH - not able to aggregate bandwidth, when applying the policy. PR1398000
- The rpd core file might be generated when Layer 2 VPN is used. PR1398685
- The rpd might crash in a BGP setup with NSR enabled. PR1398700
- UHP behavior is not supported for LDP to SR stitching scenario. PR1401214

- BGP router on the same broadcast subnet as its neighbors might cause IPv6 routing issue on the neighbor from other vendors. PR1402255
- Memory leaks when labeled IS-IS transit routes are created as chain composite next-hop. PR1404134
- Extended traffic loss might be seen after link recovery when **source-packet-routing** is used on OSPF P2P links. PR1406440
- Race conditions during BGP peer establishment causes an rpd crash. PR1410553

Services Applications

- IPsec-VPN IKE security associations might get stuck in Not Matured state. PR1369340
- Twice NAT is not supported on FTP ALG and causes an MS-PIC crash. PR1383964
- L2TP subscribers might be stuck in init state in a corner case. PR1391847
- The spd might crash when **any-ip** is configured in the **from** clause of the NAT rule with the static translation type. PR1391928
- IP ToS bits are not copied to the outer IPsec header. PR1398242
- Invalid Layer 4 checksum might be observed on IPv4 packets generated by NAT64 with MS-DPC after translating fragmented IPv6 UDP/TCP packets. PR1398542
- The ICMPv6 packet with embedded IPv6 fragment might not be translated correctly to IPv4 ICMP packet in a NAT64 with MS-DPC deployment. PR1402450
- Inconsistent content might be observed in the access line information between ICRQ and PPPoE messages. PR1404259
- The stale si logical interface might be seen when L2TP subscribers with duplicated prefixes or framed-route log in. PR1406179
- The kmd process might crash on MX and ACX platforms when IKEv2 is used. PR1408974
- The jpppd core file is seen on LNS. PR1414092
- L2TP LAC might not tunnel static PPP subscriber when you add or change interface events for related PPP logical interface that comes in a short time interval. PR1416016

Subscriber Access Management

- Address pool does not correctly cycle to the beginning of the pool when the **linked-pool-aggregation** parameter is defined. PR1374295
- The subscribers might be stuck in terminating state if RADIUS redirect is used. PR1376265
- RADIUS VSAs, Actual-Data-Rate-Downstream, and Actual-Data-Rate-Upstream values are not complaint with RFC 4679. PR1379129
- CoA updates subscriber with original **dynamic-profile** if RADIUS has returned different dynamic-profile name. PR1381230

- Some subscribers fail to get SRL service as provided in RADIUS accept message even though the RADIUS messages can be sent and received. PR1381383
- The value of **predefined-variable-defaults routing-instances** overrides the RADIUS-supplied VSA (26-1 Virtual-Router). PR1382074
- The RAA message might consist of additional AVP Destination-Host even though it is not configured for Gx-Plus session. PR1384011
- The **authd: gx-plus: logout: wrong state for request session-id <xyz>** log message is seen when a subscriber is manually Llogged out using the **clear network-access aaa subscriber username <xyz>** command. PR1384599
- Multiple IPv6 IANA addresses are assigned for one session in IPv6 PD binding failure scenarios. PR1384889
- Usage-Monitoring-Information AVP as part of PCRF Gx-plus provisioning is causing service accounting activation. PR1391411
- The DHCPv6-PD client connection might be terminated after commit when the RADIUS-assigned address is not defined within the range of a local pool. PR1401839
- An authd crash might be seen due to a memory corruption issue. PR1402012
- JSRC uses RADIUS Service accounting protocol instead of JSRC for SRC installed service. PR1403835
- The log message **authd[18454]: %DAEMON-3-LI: liPollTimerExpired returned 0** can be seen after any LI activity. PR1407923

User Interface and Configuration

- The **max-db-size** configuration do not work on some MX Series routers. PR1363048
- The **show configuration** and **rollback compare** commands are causing high CPU utilization. PR1407848

VPNs

- The receivers belonging to a routing instance might not receive multicast traffic in an Extranet next-generation MVPN scenario. PR1372613
- The **accept-remote-source** statement configured on the core interface might cause traffic outage. PR1375716
- High rpd CPU utilization on the backup Routing Engine might be observed in a MVPN with NSR scenario. PR1392792
- The rpd process crashes when the LSP template for a provider tunnel is changed. PR1395353
- Downstream interface is not removed from multicast route after getting PIM prune. PR1398458

Resolved Issues: 18.3R1

Application Layer Gateways (ALGs)

- IKEv2 negotiation might fail with the IKE ESP ALG enabled in an IKEv2 redirection scenario. PR1329611

Authentication and Access Control

- The client moves back to connecting state when VSTP is enabled along with dynamic VLAN assigned once the port gets authenticated by dot1x. PR1304397
- DHCP security is not working on MX Series platform. PR1354855
- On all Junos OS products, dynamic filter is retained if the filter attribute is not present in change of authorization (CoA). PR1364156

Class of Service (CoS)

- Remove CoS IDL from the jet IDL package and update the documentation for the same. PR1347175
- The Routing Engine might get into amnesiac mode after restarting when **excess-bandwidth-share** is configured. PR1348698
- CoS traffic control profiles might fail to apply on an aggregated Ethernet interface in a corner scenario. PR1355498
- 802.1P bit rewrite in inner-vlan header is not processed after a rewrite rule add or delete for a logical interface under the Packet Forwarding Engine. PR1375189

EVPN

- In an EVPN-VXLAN, the MAC entry is incorrectly programmed in the Packet Forwarding Engine, leading to some traffic being dropped or silently discarded. PR1231402
- MPLS label leak leads to label exhaustion and the rpd process crashes. PR1333944
- In an EVPN-VXLAN environment, BFD flap causes VTEP flap and then the Packet Forwarding Engine process crashes. PR1339084
- Traffic loss might be observed in an EVPN-VPWS scenario if the remote PE device interface comes down. PR1339217
- In EVPN-VXLAN scenarios, the traffic might get silently dropped and discarded to interfaces that are down, but LACP is up. PR1343515
- Traffic might be lost on Layer2 and Layer3 spine node in a multihomed EVPN scenario. PR1355165
- EVPN IRB configured with **no-gratuitous-arp-request** is still sending gratuitous ARP. PR1356360
- The rpd might crash if the EVPN instance refers to a vrf-export policy that does not have “then community”. PR1360437
- Proxy ARP might not work as expected in an EVPN environment. PR1368911

Forwarding and Sampling

- The LSP might take up to 30 seconds to come up when creating a policing filter and applying the filter to the LSP are both configured in a single commit. PR1160669
- DHCP service crashes after a switch or router is set to factory default by being cleared. PR1329682
- Junos OS allows firewall filters with the same name at the **[edit firewall]** and **[edit firewall family inet]** hierarchy levels. PR1344506
- The remote MAC might not be added in the forwarding table, which might cause traffic to be dropped in an EVPN scenario with RSVP and CBF configured. PR1353555
- The backup Routing Engine might write dummy interface accounting records after GRES. PR1361403

General Routing

- In timing hybrid mode, MX Series MPC2 cards are not working with ACX Series routers with VLAN (native-vlan-id). PR1076666
- The chassis alarm message **Bottom Fan Tray Pred Fail** needs to be rewritten so that the meaning is less obscure. PR1202724
- Tacacs access does not work after upgrade. PR1220671
- An incorrect TBB Packet Forwarding Engine component temperature might be reported on the MX80. PR1259379
- On MX Series, the **show chassis led** command should not be displayed in possible completions of the **show chassis** command. PR1268848
- Flexible PIC concentrator (FPC) crash/reboot is observed when bringing up about 12,000 Layer 2 bit stream access (L2BSA) subscribers simultaneously. PR1273353
- Error messages are observed on the vty session while running a script for IGMP snooping over EVPN-VXLAN. PR1276947
- Migration from syslog API to errmsg API **/src/junos/usr.sbin/mobiled** is observed. PR1284625
- Migration from syslog API to errmsg API **/src/junos/usr.sbin/mspmand** is observed.PR1284643
- In an EVPN-VXLAN interface scenario, inter-vrf traffic black hole occurs after repeated restart of routing on redundant gateways. PR1289091
- PPPoE cannot dial in due to PADI being dropped as **unknown iif** when the aggregated Ethernet interface configuration is deactivated or activated. PR1291515
- SSH to the Ubuntu-based JDM is not stable. PR1291836
- The rpd might crash by executing the command **show route extensive** during deletion of the IS-IS configuration. PR1301849
- Incorrect packet statistics is reported in ifHCInUcastPkts OID. PR1306656

- The error messaged **pfeman_inline_ka_steering_gencfg_handler** might be seen during FPC restart with BFD configured. PR1308884
- Subscribers might not be able to access the device if dynamic VLAN is used. PR1309770
- On the MX10000 need to suppress the chassis alarm for switched-off PEMS. PR1311574
- The L2TP LAC might drop packets that have incorrect payload length while sending packets to the LNS. PR1315009
- CoS is not applied to the Packet Forwarding Engine when a VCP link is added. PR1321184
- The rpd might crash when two next hops are installed with the same next-hop index. PR1322535
- The CLI command **request vmhost halt routing-engine other** does not halt the backup Routing Engine. PR1323546
- Migration from syslog API to errmsg API **/src/junos/usr.sbin/aaad** is observed. PR1327266
- With auto-installation USB configured, interface-related commits might not take effect due to dcd error. PR1327384
- When an AMS bundle has a single mams-interface added to it, the subinterfaces do not recover after the subinterface has been disabled. PR1329498
- The host outbound traffic cannot be rewritten for IEEE-801.1p bit in a dynamic subscriber over PS interface scenario. PR1329555
- The **Too many supplies missing in Lower/Upper zone** alarm flaps (set/clear) every 20 seconds if a zone does not have the minimum number of required PSMs. PR1330720
- Juniper Development Innovation Diagnostics (JDID) thrashes continuously and continuous log messages are observed in syslog. PR1333632
- Two subscribers cannot reach the online state at the same time if they have an identical frame-route attribute value. PR1334311
- Tc_count counters in a filter with the **scale-optimized** statement, are not incrementing. PR1334580
- MPC5E line cards went for "restart" after a unified ISSU to Junos OS Release 18.2DCB in MX2010 box. PR1334612
- The master LED glows on the master and backup Routing Control Board during an image upgrade on the master with GRES/NSR enabled. PR1335514
- The RIP route updates might be partially dropped when NSR is enabled. PR1335646
- On MPC7E, ukern crashes and FPC reboots with vty command **show agent sensors verbose**. PR1366249
- MS-MPC/MS-PIC might crash in NAT scenario. PR1366259
- With certificate hierarchy, where intermediate CA profiles are not present on the device, in some corner cases, the pkid can become busy and stop responding. PR1336733
- The hash value generated for 256-bit key length of AES-GCM-256 algorithm is incorrect PR1336834

- AI-script can be manually upgraded after a Junos OS upgrade. PR1337028
- Links are flapping or staying down due to an interoperability issue between the MX Series router (or EX9200 switch) and the transport device. PR1337327
- MPC throughput degradation might be seen after SBF2 goes online or offline. PR1338216
- CLI shows CB states online after you press the RCB offline button for more than 4 seconds. PR1340431
- A few subscribers show the wrong accounting values in a large-scale subscribers scenario. PR1340512
- VRRP gets stuck on the master during upgrade or cold boot. PR1341044
- IPv4 or IPV6 traffic is routed out through the wrong interface after rpd restarts the leaf device in the IP-CLOS profile. PR1341381
- Reboot of the Routing Engine might occur if the PPPoE interface is configured over an aggregated Ethernet or RETH interface. PR1341968
- SNMP walk might fail for LLDP-related OIDs. PR1342741
- The vFPC might become absent, resulting in the total loss of traffic. PR1343170
- In an MPLS or RSVP environment, LSP might get stuck in DN state with **Record route: <self> ...incomplete**. PR1343289
- On upgrading from Junos OS Release 18.1 to Junos OS Release 18.2 DCB image, errors are observed in a unified ISSU because of the ffp process. PR1343542
- MPC8/9E card crashes and generates a core file during logout of DHCPv6 subscribers over on static VLAN. PR1343965
- The RLT interface might not be able to route and forward traffic in Junos OS Release 17.3. PR1344503
- The framed-route "0.0.0.0/0" cannot be installed on MX Series platforms with Junos OS enhanced subscriber management releases. PR1344988
- The ARP reply packet automatically generates the virtual gateway MAC address in the Ethernet header. PR1344990
- In a Junos Fusion Enterprise, there is an issue with 802.1X reauthentication. PR1345365
- An rpd crash might be seen if the **no-propagate-ttl** statement is set in a routing instance that has a specific route. PR1345477
- The Routing Engine model is changed from JNP10003-RE1 to RE-S-1600x8. PR1346054
- Additional **show** commands are called when the **request support information** command is issued. PR1346129
- New PPPoE users might fail to log in. PR1346226
- **AC system error** counter in **show pppoe statistics** is not working. PR1346231
- VCCP-ADJDOWN detection is delayed on the Virtual Chassis backup router when one VCP link is deleted on the Virtual Chassis master router. PR1346328

- The twice-napt-44 sessions are not synchronizing to the backup SDG with stateful synchronize configured. PR1347086
- IPv6 MAC address resolution might fail if the DHCPv6 client uses a non-EUI64 link-local address. PR1347173
- The Juniper Network devices running Junos OS might encounter a chassis alarm indicating **FPC 0 Major Errors - PE Error code: 0x2100ba**. PR1347805
- The rpd might crash when the dynamic tunnels next hop resolving migrates to a more specific IGP route. PR1348027
- The FPC might crash because of the MIC error interrupt hogging. PR1348107
- Packet loop is detected in the EIBGP multipath environment with an **install-nexthop** policy enabled. PR1348175
- Unable to set fti as output for **port-mirroring** instance. PR1348317
- Chassisd memory leak issue is observed on MX10003 and MX204 platforms, and it might eventually cause chassisd crash and Routing Engine switchover. PR1348753
- In certain scenarios on MX Series Virtual Chassis with L2TP LNS, the DHCPv6 solicit packet might be dropped. PR1348846
- Routing Engine mastership keepalive timer is not updated after the GRES configuration is removed. PR1349049
- The dcd process might crash after any other smid-related daemon crashes. PR1349154
- The major alarm **Major PEM 0 Input Failure** might be observed for DC PEM. PR1349179
- The mspmand process might crash when executing the **show services nat deterministic-nat nat-port-block** command. PR1349228
- Mgd crashes and generates a core file because of an issue in nsindb infrastructure. PR1349288
- When VOIP VLAN is set as NATIVE VLAN on the port, the interface still shows up as a tagged interface and drops all untagged traffic. PR1349712
- PS over rLT does not work on MPC7, MPC9; PS over LT for the same scenario works. PR1350115
- The pccd might crash after a delegated LSP is removed in a PCEP scenario. PR1350240
- Stale access-internal routes corresponding to BOUND interfaces (clients) might remain in rpd when AIU temporarily fails before succeeding eventually. PR1350401
- The MTU value for the subscriber's interface might be programmed incorrectly if the statement **routing-services** or **protocol pim** is configured in a dynamic profile. PR1350535
- The VCP port might not come back up after it is removed and added again. PR1350845
- The subinfo process might crash when the executing **show subscribers address <> extensive** command for a DHCPv6 address. PR1350883

- PPE **asynchronize extension error** occurs when FPC is restarted or removed. PR1350909
- The pfed process might consume high CPU resources if subscriber or interface statistics are used at a large scale. PR1351203
- Dynamic physical interface creation fails when the SFP optic is plugged in the MX150. PR1351387
- High CPU usage of the bbe-smgd process might be seen when L2BSA subscribers get stuck. PR1351696
- After GRES, the BGP neighbors at the master Routing Engine might reset, and the BGP neighbors at the backup Routing Engine take a long time to establish. PR1351705
- Multicast route might flap when ephemeral database is enabled. PR1352499
- Junos node slicing MSE after reinstall causes one JDM server to complains. The pull configuration fails and the system falls back to the push configuration method. PR1352503
- The bbe-smgd daemon might restart in a subscriber environment. PR1352546
- The DHCP relay-reply packets are dropped in the DHCPv6 relay scenario. PR1352613
- CM error CLI is not working on the Junos Node Slicing. PR1352705
- Taking the MIC6-100G-CFP2 MIC offline by using the CLI command might trigger an FPC card crash. PR1352921
- Migration from syslog API to errmsg API **/bbe-svcs/smd/plugins/cos/** is observed. PR1353179
- The rpd is permanently hogging CPU resources due to a logical system configuration commit. PR1353548
- The 3D 40x 1GE(LAN) RJ45 MIC is not recognized on the MX104. PR1353632
- Traffic interruption is observed after multiple Routing Engine switchovers. PR1354002
- Observing chasissd crash after chassisd restart in MX10003. PR1354269
- The syslog error **dfw_bbe_filter_bind:1125 BBE Filter bind type 0x84 index 167806251 returned 1** is observed. PR1354435
- The rpd process crashes and generates a core file when adding an inter-region template in routing instances. PR1354629
- Aggregated Ethernet operational state goes up even though some of the member interfaces configured under the aggregated Ethernet interface are down. PR1354686
- The ifinfo process could crash on MX Series routers with BNG running L2BSA service. PR1354712
- The static-subscribers do not properly update firewall information on the Packet Forwarding Engine when dynamic configuration changes are made to active subscribers. PR1354774
- There is memory leak on agentd when Junos Telemetry Interface is configured. PR1354922
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100 Gbps). PR1355168
- Alarm LED is not working in MX204 to indicate the minor or major faults. PR1355225

- Packets destined to the Routing Engine might be dropped in the kernel when LACP is configured. PR1355299
- Syslog message is observed during a unified ISSU. PR1355345
- Fabric chip failure alarms are observed in a GRES scenario. PR1355463
- Syslog messages **ui_client_connect_to_kmd_instance: KMD-SHOW connect to kmd-instance failed kmd-instance Routing Engine, fpc slot 0, pic slot 0** are seen. PR1355547
- The chassis alarm is not reflecting the correct state when INP0 and INP1 have out-of-range AC voltage. PR1355803
- The **flex-flow-sizing** is not working on the MX204. PR1356072
- The MPLS IPv4 templates do not have correct src AS and dst AS as 4294967295, and src Mask and DstMask as 0 after adding mpls-flow table size occurs on the fly. PR1356118
- The rpd process crashes when issuing the command **show dynamic-tunnels database terse** for RSVP automatic mesh tunnels. PR1356254
- L2c messages from PEM and PSM are reported if SNMP is enabled. PR1356259
- Executing the command **show pppoe underlying-interfaces** might cause the bbe-smgd to crash in a scaling subscriber environment. PR1356428
- Link stays up unexpectedly on MX204 with copper cable removed. PR1356507
- DHCP subscribers fail after reconfiguration of port from tagged to untagged mode. PR1356980
- Starting with Junos OS Release 18.2R1, PTPoE packet exchanges do not happen with the MIC-3D-SR-4GE-2XGE when PTP master and slave interfaces have "ethernet-bridge" encapsulation and are part of a bridge domain. PR1357017
- The bbe-smgd process might get stuck in subscriber scenario with node slicing. PR1357252
- Upgrading from Junos OS Release 15.1F2-S20 to Junos OS Release 15.1X12 using "validate" results in a **Fabric Mixed Mode error**. PR1357423
- Routing Engine switchover that occurs before the backup Routing Engine is GRES ready might cause line card restart, Routing Engine kernel crash and multiple chassisd crashes. PR1357427
- The rpd memory leak occurs with RT_NEXTHOPS_TEMPLATE. PR1357897
- Traffic might be sent to an incorrect RLT member interface after RLT switchover. PR1358320
- Incorrect traffic load balance might be seen even if **locality-bias** is configured on the MX Series Virtual Chassis. PR1358635
- The **show chassis fpc** command output might show "Bad Voltage" for an FPC powered off by the configuration or the CLI command after the command **show chassis environment fpc** is executed. PR1358874

- The bbe-smgd process might crash and generate a core file at #6 0x00000000006937ad in bbe_set_index (type=<optimized-out>, bbe_index= <optimized out>) at **../../../../../../src/junos/usr.sbin/bbe-svcs/smd/infra/bbe_index.c:459**. PR1359290
- FRU-model-number is not displayed for few FRUs in /component sensor for the MX10008 and MX10003 platforms. PR1359300
- The IPv6 subscriber might fail to access the network. PR1359520
- The PSTP subscriber might not be able to log in on the BNG device. PR1359574
- During a scheduled boot, both Routing Engines might fail with a special time format. PR1359602
- PluginExit() function is never called. PR1359610
- Bbe-smgd might fail to add members to some of the aggregated Ethernet interfaces at random when there are many aggregated Ethernet in the access configuration. PR1359986
- The rpd crashes and generates a core file at **../../../../../../../../src/junos/usr.sbin/rpd/lib/rt/rt_attrib.c**, line 3329: "rt_template_get_rtn_ngw(nhp) <= 1" ' on doing Routing Engine switchover with SRTE routes. PR1360354
- FPC core file is observed after GRES switchover in RE1 at sensor_export_get_format. PR1361015
- The rpd scheduler slip might be seen when frequently deleting, modifying, or adding groups that are applied on the top level. PR1361304
- The rpd process get struck at 100 percent after clear bgp neighbor operation. PR1361550
- Migration from syslog API to errmsg API **usp/usr.sbin/nsd/common/nsd_tpm.c** is observed. PR1361986
- Spontaneous bbe-smgd process might generate a core file on the backup Routing Engine. PR1362188
- Executing **show route prefix proto ip detail** during route churm in a route scale scenario might lead to FPC crash. PR1362578
- Unexpected DCD_PARSE_ERROR_SCHEDULER messages are logged when MS-MPC and MS-MIC are brought offline or online. PR1362734
- A quick memory leak might be seen in the bbe-smgd daemon if the dynamic profile variable name and the default associated value are configured to be the same. PR1362810
- The non-default routing-instance is not supported correctly for NTP packet in a subscriber scenario. PR1363034
- Traffic destined to the MAC/IP address of VRRP VIP gets dropped on the platforms that have common TFEB terminals such as MX5, MX10, MX40, MX80, and MX104. PR1363492
- The pmbus_read_volt: sfb-07 - MAX20751-PF1-0.9v: pmbus read failed for cmd 0x8b. PR1363587
- The xmlproxyd for internal interfaces is reporting uint32 instead of uint64. PR1363766
- The l2circuit on MPC7E, MPC8E, MPC9E with asynchronous-notification and ccc configured might keep flapping when the circuit is going up. PR1363773

- A traffic loop might occur even though that port is blocked by RSTP in a ring topology. PR1364406
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with the link-layer-down flag set. PR1365263
- Traffic is dropped when a three-link training failure is seen in a line card. PR1365668
- An upgrade to Junos OS Release 18.1R1.9 fails. Installing package nfx-2-routing-data-plane-1.0-0.x86_64 needs 76 MB on the / filesystem. PR1366324
- Migration from syslog API to errmsg API **junos/lib/liboiu-ffp/** is observed. PR1366546
- The next hop of the MPLS path might get stuck in hold state, which might cause traffic loss. PR1366562
- SNMP MIB walk for UDP flood gives different output statistics than the CLI. PR1366768
- Taking the fabric links of PFE 4 and PFE 5 offline is not supported. PR1367412
- The bbe-smgd crashes if an L2BSA subscriber receives a routing instance name where VPLS is not configured. PR1367472
- The **show system resource-monitor fpc** command might show a nonexisting Packet Forwarding Engine. PR1367534
- RTG interface status will be shown as incorrect status with show interfaces. PR1368006
- The authd process might not be started after executing Routing Engine switchover on the backup Routing Engine without GRES enabled. PR1368067
- Multiple provisioning and deprovisioning cycles cause rdmd memory leak. PR1368275
- For a route resolved to a next hop with multiple gateways, and some of the gateways were rejected during the route resolution, then the final next-hop result might contain incorrect gateway formation. RPD API rt_nexthops_extract_gateway_convert_unnumbered_gf_dli() rectification. PR1368855
- The **commit** or **commit check** command might fail because of the error **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- The subscriber filter is not removed from the Packet Forwarding Engine when routing-services are enabled in the dynamic profile on an L2TP LNS. PR1369968
- Kernel crash might be seen after committing a demux-related configuration. PR1370015
- The rpd might crash after Routing Engine switchover is performed or the rpd is restarted if interface-based Dynamic GRE Tunnel is configured. PR1370174
- Packets that exceed 8,000 bytes might be dropped by MS-MPC in an ALG scenario. PR1370582
- SFP-1FE-FX optics is not coming up on GMIC. PR1370962
- All the MX150 devices running VRRP on a LAN are stuck in master state. PR1371838
- FPC high CPU utilization or crash occurs during a hot-banking condition. PR1372193

- The smgd process crashes and generates a core file after essmd restarts with reference to mmf_ensure_mapped (mmf=0xe8f0200, offset=4294967295, len=108) at **../src/junos/lib/libmmf/mmf.c:1972**. PR1372223
- On a high scale l3vpn, traffic is dropped when egressing on an AF interface. PR1372310
- The Routing Engine might crash after a non-GRES switchover. PR1373079
- BOOTP packets might get dropped if BOOTP-support is not enabled at the global level. PR1373807
- A vMX QoS performance issue occur in Junos OS Release 18.3. PR1373999
- The bbe-smgd crashes and generates a core file continuously while deleting a multicast group node from the tree. PR1374530
- PCE-initiated LSPs remain "Control status became local" after removing PCE configuration. PR1374596
- The rpd generates core files at **io_session_trace ioth_read_request_process jtask_jthr_thread_main_loop**. PR1374759
- The ICMPv6 packets larger than 1024 might be dropped if **icmp-large-packet-check** is configured on IDs service. PR1378852

High Availability (HA) and Resiliency

- The backup Routing Engine might go to db prompt after performing a configuration remove and restore. PR1269383
- The ksyncd process might crash continuously on the new backup Routing Engine after performing GRES. PR1329276
- The Virtual Chassis backup router cannot synchronize with the Virtual Chassis master router when the Virtual Chassis splits then reforms. PR1361617

Infrastructure

- Cleanup at thread exit causes memory leaks. PR1328273
- The fxp0 interface does not accept IP address with **master-only** applied. PR1341325
- The kernel might crash and the system might reboot in a SNMP query reply scenario. PR1351568
- Junos OS is no longer going to db prompt at ~ + ctl-b. PR1352217

Interfaces and Chassis

- L2TP subscribers might not be cleared if the access-internal routes fail to install PR1298160
- Subscribers might fail to access the device after deleting the needless aggregated Ethernet configuration. PR1322678
- When in hardware-assisted-pm-mode and pm configuration is scale, deactivating eth-oam can lead to an FPC crash. PR1347250
- Suppressing cfmd logs: jnxSoamLmDmCfgTable_next_lookup: md 0 ma 0 md_cfg 0x0. PR1347650

- The PPPoE subscribers might fail to login for authd running on 100 percent CPU with high frequency of On-Demand IP address allocation requests. PR1348578
- Spontaneous jpppd core file is generated on the backup Routing Engine in longevity test at **../../../../../src/junos/usr.sbin/jpppd/pppMain.cc:400**. PR1350563
- VRRP VIP becomes unreachable after deleting one of the logical interfaces PR1352741
- **native-vlan-id** support on ps-interface. PR1352933
- The FPC might be stuck at 100% for a long time when MC-aggregate Ethernet with enhanced-convergence is configured with large-scale IFLs. PR1353397
- Any filter change applied to a FTI interface triggers the FTI interface flap PR1354832
- The aggregate Ethernet interface might flap when the link speed of the aggregate Ethernet bundle is configured to oc192 PR1355270
- FPC core related to cfmman PR1358192
- Clients might not get IPv4 addresses in a PPPoE dual-stack scenario. PR1360846
- Approximately 50 percent of PPPoE subscribers (PTA and L2TP) and all ESSM subscribers are lost after ISSU during DT CST stress test PR1360870
- Starting with Junos of Release 17.2R1, the CLI allows you to configure more than 2048 logical interfaces on the LAG interface. PR1361689
- Error messages like **ifname [ds-5/0/2:4:1] is chan ci candidate** are seen during a commit operation. PR1363536
- The EOAM LTM messages might not get forwarded after system reboot in a CFM scenario configured with the CCC interface. PR1369085
- Subscribers cannot negotiate an MLPPP session with MX Series LNS when the dynamic-profile name contains more than 30 characters. PR1370610
- The dcd process might go down when **vlan-id none** is configured for the interface. PR1374933
- FTI logical interface VNI limits changed from (0..16777215) to (0..16777214) PR1376011
- Duplicate IP cannot be configured on both sonet (so-) interface and other interfaces. PR1377690

Layer 2 Ethernet Services

- The MAC address might not be learned due to spanning-tree state "discarding" in kernel table after Routing Engine switchover. PR1205373
- Migration from syslog API to errmsg API **/src/junos/usr.sbin/lacpd** is observed. PR1284592
- The DHCPv6 second Solicit message might not be processed when IA_NA and IA_PD are sent in a separate Solicit message. PR1340614
- The DHCP client is not able to connect if VLAN is modified on the aggregated Ethernet interface associated with the IRB interface. PR1347115

- When DHCP subscribers are in BOUND (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state, if dhcp-service is restarted, then the subscribers in this state are logged out. PR1350710
- The DHCP relay agent will discard the DHCP request message silently if the requested IP address has been allocated to the other client. PR1353471
- Restarting the FPC that hosts the micro-bfd link might cause LACP to generate a core file. PR1353597
- DHCPv6 relay ignores replies from the server when renewing. PR1354212
- Migration from syslog API to errmsg API PPMD client LACP is observed. PR1358599
- The DHCP leasequery message is replied to with an incorrect source address. PR1367485
- A rebinding state counter is added to DHCPv4 and DHCPv6 binding sensors. PR1368392

Layer 2 Features

- VPLS instance stays in NP state after LDP session flaps. PR1354784
- The Routing Engine kernel might crash when OSPFv3 is configured with IPsec key authentication over the IRB interface. PR1357430
- The dcpfe/fxpc process might crash on Packet Forwarding Engines with low memory when allocating huge memory. PR1362332
- The traffic might not be transmitted correctly in large scale of VPLS scenario. PR1371994

MPLS

- When the explicit fate-sharing group cost is removed from the configuration, the default value "1" should be used in calculations. PR1330161
- After a MPLS LSP link flap and local repair, RSVP tries to create a new LSP instance, but the instance might get stuck. PR1338559
- An rpd crash might happen in an RSVP setup-protection scenario. PR1349036
- Some LSPs might be stuck on the upstream devices after interfaces flap occurs on downstream devices. PR1349157
- In a very rare scenario, rpd might crash when LDP failed to allocate self-id for the P2MP FEC. PR1349224
- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. PR1359956
- L2 circuit might flap after an interface goes down even if the LDP session stays up when l2-smart-policy is configured. PR1360255
- The rpd process might crash during P2MP LSPs churn. PR1363408
- The rpd might crash in a BGP LU and LDP scenario. PR1366920
- The traceroute MPLS LDP to a Huawei fails until TTL expires. PR1372924

- The traffic might not be load-balanced equally across LSPs with **ldp-tunneling** configured. PR1373575
- The rpd process might crash continuously if nsr-synchronization or all flag is used in RSVP traceoptions. PR1376354

Multicast

- Some IGMP groups might have the wrong upstream interface because the discard route is installed in the PIM. PR1337591

Network Management and Monitoring

- Output for the **show pfe statistics traffic** command output shows traffic statistics as zero for a brief time after doing "test panic" on a non-traffic-carrying line card. PR1349517
- EVENTD fails to start up with syslog configuration. PR1353364
- The jnxDcuStatsEntry and jnxScuStatsEntry OIDs are missing after interface configuration changes. PR1354060
- SNMP process crashes when polling CFM statistics. PR1364001

Platform and Infrastructure

- The command **show configuration | compare** shows the unchanged configuration after deleting part of the configuration. PR1042512
- Error messages might be observed with MPC5E card. PR1283850
- The apply-path prefix is not inherited under the policy after commit. PR1286987
- Need to move **XQ_CMERROR_XR_CORRECTABLE_ECC_ERR** to minor and reclassify remaining **XQCHIP CMERROR** from fatal to major. PR1320585
- On the MX104, the backup Routing Engine kernel crashes on committing **set system management-instance**. PR1335903
- Configuring the same DHCP server in different routing instances is not supported in DHCP relay scenario. PR1342019
- The interface remains down after **delete interface <int> disable**. PR1343317
- ZTP is not supported for vmhost images on next-generation Routing Engines on MX Series platforms. PR1343338
- On the MPC5, inline-ka PPP echo requests not transmitted when the anchor-point is lt-x/2/x or lt-x/3/x in a pseudowire deployment. PR1345727
- Multiple vulnerabilities exist in cURL. For more information, refer to JSA10874. PR1347361
- The IPv4 GPRS traffic over aggregated Ethernet interface might be dropped if gtp-tunnel-endpoint-identifier is configured. PR1347435
- EVPN-VXLAN, MX Series: Output policing action does not work on IRB interfaces for VNIs. PR1348089

- FPC CPU utilization with LT interfaces is pegged continuously at 100 percent. PR1348840
- Running RSI through the console port might cause the system to crash and reboot. PR1349332
- ICMP error messages are not generated if “don't fragment” packets exceed the MTU of the multiservice interface. PR1349503
- **[ui]** Some commands of **system ddos-protection protocols unclassified** are missing on MX2020 in Junos OS Release 17.2X75. PR1349782
- When viewing IPv6 addresses, **display rfc5952** does not work when combined with **display set**.PR1349949
- The lt- interface gets deleted with the tunnel-services configuration still present. PR1350733
- Chassis manager daemon (chassisd) memory leak occurs. PR1353111
- In a Junos Fusion setup, configuring VRRP on an extended port will lead to a kernel crash. PR1353498
- The FPC would crash due to the memory leak caused by the VTEP traffic. PR1356279
- A traffic black hole is seen along with the message **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH messages**. PR1357707
- When the **forwarding-class-accounting** statement is enabled on an interface, inside of a routing-instance of instance-type vrf, aggregate input forwarding-class statistics do not increment (egress statistics work fine). PR1357965
- Select CLI functions are not triggering properly (for example, **set security ssh-known-hosts load-key-file** and **set system master-password**). PR1363475
- Authentication for adding the DTCP filter is not happening on the router and the filter is not getting added. PR1365515
- The same vlan-id is not allowed on multiple logical interfaces of the same GR interface. PR1365640
- Qmon Sensors not working with hyper-mode enabled. PR1365990
- Subscribers over aggregated Ethernet interface might have tail drops, which will affect the fragmented packets due to the QXCHIP buffer getting filled up. PR1368414
- The host outbound traffic might get dropped when the **class-of-service host-outbound-traffic ieee-802.1 rewrite-rules** statement is configured. PR1371304
- The logical tunnel interface might be unable to send out control packets generated by the Routing Engine. PR1372738
- JNH memory leaks occur in multicast scenario with MoFRR enabled. PR1373631

Routing Policy and Firewall Filters

- The policy might not get cleaned up after a configuration is deleted, which could cause an rpd to generate a core file. PR1357724

Routing Protocols

- BGP extended communities with sub-type 4 are erroneously displayed at LINK_BANDWIDTH. PR1216696
- The rpd generates core files in the ASBR when BGP is deactivated in the ASBR before all stale labels have been cleaned up. PR1233893
- Migration from syslog API to errmsg API **/src/junos/usr.sbin/ppmd** is observed. PR1284621
- Multihop eBGP peering session exchanging EVPN routes can result in an rpd core file when BGP updates are sent. PR1304639
- The primary path of MPLS LSP might switch to another address. PR1316861
- The mcsnoopd process memory leak occurs. PR1326410
- OSPF rLFA default PQ node selection algorithm does not provide proper protection paths in a large-scale network. PR1335570
- Changes to the displayed value of AIGP occur with the **show route ... extensive** command. PR1342139
- A traffic black hole might be seen if the local device is receiving BFD-down. PR1342328
- The resetting of SRTE sensors is not predictable after the rpd is restarted (restart-routing). Transit sensors are reset all the time but ingress sensor resetting is unpredictable. PR1345229
- The rpd process might crash after GRES when multipath is configured. PR1346954
- The rpd might generate a core file when running streaming telemetry. PR1347431
- The rpd might crash if a route for RPF uses a qualified-next-hop. PR1348550
- The rpd might crash while restarting routing or deactivating IS-IS. PR1348607
- The rpd might crash after executing Routing Engine switchover. PR1349167
- Traffic loss might be seen after the upstream interface shifts from one to another when receiving the PIM prune packet. PR1350806
- The rpd might crash when BGP route damping and BGP multipath feature are configured. PR1350941
- The **source-as community** statement is not appended to RP (display issue in **show route** detail output). PR1353210
- Static route flap occurs on commit when configured with resolve statement. PR1366940
- On MX Series Virtual Chassis, a 10 minute traffic loss might be caused by BGP flap during a unified ISSU. PR1368805
- Route entry might be missing when IS-IS shortcut is enabled and MPLS link flap. PR1372937

Services Applications

- SNMP MIBs are not yielding data related to sp- interfaces. PR1318339
- The software should selectively start the ZLB delay timer at the Packet Forwarding Engine for LAC tunnels. PR1338450
- The bbe-smgd process might crash if there are 65,535 L2TP sessions in a single L2TP tunnel. PR1346715
- Session limit per tunnel on LAC does not work as expected. PR1348589
- When performing an SNMP walk on the IKE SA that is deleted, IPsec tunnels might go down and an infinite loop scenario might be seen. PR1348797
- UDP checksum inserted by MS-DPC after NAT64 is not valid when incoming IPv4 packet has UDP checksum set to 0. PR1350375
- The **show services stateful-firewall flows** counter shows exceedingly high numbers. PR1351295
- Jl2tpd process might crash shortly after one of the L2TP destinations becomes unavailable. PR1352716
- IPsec tunnels might flap when SNMP walk is executed if IPsec is configured with DPD enabled. PR1353240
- In an L2TP, tunnel-switch clients in the subscriber session database reference the incorrect routing instance. PR1355396
- L2TP access concentrator (LAC) tunnel connection request packets might be discarded on LNS device. PR1362542
- Some subscriber might be stuck in terminating state in L2TP scenario. PR1363194
- The L2TP subscribers might not be able to log in successfully because of the jl2tpd memory leak. PR1364774
- Accounting stop message is not sent to RADIUS server after bringing down the L2TP subscriber. PR1368840
- Actual data rate downstream value is not included in the L2TP ICRQ message from the LAC. PR1370699
- NAT64 does not translate ICMPv6 Type 2 packet (packet is too big) correctly when MS-DPC is used for NAT64. PR1374255

Subscriber Access Management

- Multiple RADIUS servers having different dynamic request ports is not supported. PR1330802
- Subscriber might get stuck in terminated state when the JSRC synchronization state is stuck in "FULL-SYNC in progress". PR1337729
- In a dual-stack subscribers scenario with NDRA pool configured, the linked pools are not used when the first NDRA pool is exhausted. PR1351765

- When attempting to scale clients, sdbsts_lock_holder.bbe-smgd.pid10686.core crashes and generates a core file. PR1358339
- CoA updates subscriber with original dynamic-profile if RADIUS has returned different dynamic-profile name. PR1381230

User Interface and Configuration

- The mustd process crashes and generates a core file ppool_bkt (phdr=0xde918024, pfile=0xde933004, no_pages=1) at **../../../../../../src/ui/lib/memory/page_pool.c**. PR1309074
- Automatic completion of interface range with ae1+TAB results in an **invalid value** error. PR1353741

VPNs

- The multicast route might be rejected when Junos OS PE devices receive C-Mcast route from other vendors' PE device. PR1327439
- The rpd crashes after committing interface-related parameters (for example, MTU change, VRF RD/RT, QoS) on PS interface with vlan-ccc encapsulation and no vlan-id. PR1329880
- The rpd might continuously crash on the backup Routing Engine and some protocols might flap on the master Routing Engine when **hot-standby** is configured for l2circuit or VPLS backup neighbor. PR1340474
- The rpd might crash on the backup Routing Engine when changing the l2circuit virtual-circuit-id in an NSR scenario. PR1345949
- The rpd process might crash after a configuration change in an L2VPN scenario. PR1351386
- In a dual-homed next-generation MVPN, the receipt of Type 5 withdrawal removes downstream join states for some routes. PR1368788

SEE ALSO

New and Changed Features | **98**

Changes in Behavior and Syntax | **116**

Known Behavior | **124**

Known Issues | **131**

Documentation Updates | **181**

Migration, Upgrade, and Downgrade Instructions | **182**

Product Compatibility | **189**

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 18.3R2 documentation for MX Series.

Subscriber Management Access Network Guide

- The *Broadband Subscriber Access Protocols User Guide* has been updated to clearly describe when you must commit changes during the process of moving the anchor point for a pseudowire subscriber logical interface device on MX Series routers. See Changing the Anchor Point for a Pseudowire Subscriber Logical Interface Device for information about moving the anchor point from one logical tunnel to another logical tunnel, from a logical tunnel to a redundant logical tunnel, and from a redundant logical tunnel to a logical tunnel.
- The guide failed to include a feature that enables you to override the information that the LAC sends to the LNS in L2TP Calling Number AVP 22 when the LAC is configured to use the Calling-Station-ID format. You can configure the access profile to override that value for AVP 22 with any combination of the agent circuit identifier and the agent remote identifier received by the LAC in the PADR packet.

  [See Override the Calling-Station-ID Format for the Calling Number AVP].
- The guide incorrectly stated that the **linked-pool-aggregation** statement is located at the **[edit access address-assignment pool *pool-name*]** hierarchy level. In fact, this statement is located at the **[edit access]** hierarchy level.

  See Configuring Address-Assignment Pool Linking.

Subscriber Management Provisioning Guide

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

Subscriber Management VLANs Interfaces Guide

- The *Broadband Subscriber VLANs and Interfaces User Guide* did not clearly indicate that only demux0 is supported for demux interfaces. If you configure a different demux interface, such as demux1, the configuration commit fails.

SEE ALSO

New and Changed Features | **98**

Changes in Behavior and Syntax | **116**

Known Behavior | **124**

Known Issues | **131**

Resolved Issues | **148**

Migration, Upgrade, and Downgrade Instructions | **182**

Product Compatibility | **189**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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 might take several minutes, depending on the size and configuration of the network.

Starting in Junos OS 18.3R2 release, FreeBSD 11.x is the underlying OS for all Junos OS platforms which were previously running on FreeBSD 10.x based Junos OS. FreeBSD 11.x does not introduce any new Junos OS related modifications or features but is the latest version of FreeBSD.

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

Platform	FreeBSD 6.x-based Junos OS	FreeBSD 11.x-based Junos OS
MX5,MX10, MX40,MX80, MX104	YES	NO
MX240, MX480, MX960, MX2010, MX2020	NO	YES

Basic Procedure for Upgrading to Release 18.3

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 Installation and Upgrade Guide.

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

Procedure to Upgrade to FreeBSD 11.x based Junos OS

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

To download and install FreeBSD 11.x based Junos OS:

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

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

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

- For 32-bit Routing Engine version:

  `user@host>` **request system software add no-validate reboot *source*/junos-install-mx-x86-32-18.3R2.9-signed.tgz**

- For 64-bit Routing Engine version:

  `user@host>` **request system software add no-validate reboot *source*/junos-install-mx-x86-64-18.3R2.9-signed.tgz**

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

- For 32-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-32-18.3R2.x-limited.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-18.3R2.9-limited.tgz
  ```

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

- ***/pathname***—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname/pathname***
  - **http://*hostname/pathname***
  - **scp://*hostname/pathname***

Do not use the **validate** option while upgrading from Junos OS (FreeBSD 6.x) to Junos OS (FreeBSD 11.x). This is because programs in the **junos-upgrade-x** package are built based on FreeBSD 11.x, and Junos OS (FreeBSD 6.x) would not be able to run these programs. You must run the **no-validate** option. The **no-validate** statement disables the validation procedure and allows you to use an import policy instead.

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

Rebooting occurs only if the upgrade is successful.

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

NOTE: After you install a Junos OS Release 18.3 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

Procedure to Upgrade to FreeBSD 6.x based Junos OS

Products impacted: MX5, MX10, MX40, MX80, MX104.

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

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

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

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

  `user@host>` **request system software add validate reboot *source*/jinstall-ppc-18.3R2.9-signed.tgz**

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

  `user@host>` **request system software add validate reboot *source*/jinstall-ppc-18.3R2.9-limited-signed.tgz**

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

- ***/pathname***—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname*/*pathname***
  - **http://*hostname*/*pathname***
  - **scp://*hostname*/*pathname***

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

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

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 18.3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

Upgrading a Router with Redundant Routing Engines

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

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

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

Downgrading from Release 18.3

To downgrade from Release 18.3 to another supported release, follow the procedure for upgrading, but replace the 18.3 jinstall package with one that corresponds to the appropriate release.

NOTE: You cannot downgrade more than three releases.

For more information, see the Installation and Upgrade Guide.

SEE ALSO

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os

New and Changed Features

IN THIS SECTION

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

Release 18.3R2 New and Changed Features

There are no new features in Junos OS Release 18.3R2 for NFX Series devices.

Release 18.3R1 New and Changed Features

There are no new features in Junos OS Release 18.3R1 for NFX Series devices.

SEE ALSO

- Changes in Behavior and Syntax | 192
- Known Behavior | 193
- Known Issues | 194
- Resolved Issues | 195
- Documentation Updates | 196
- Migration, Upgrade, and Downgrade Instructions | 197
- Product Compatibility | 199

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for NFX Series devices.

Release 18.3R2 Changes in Behavior and Syntax

There are no changes in behavior and syntax for NFX Series in Junos OS Release 18.3R2.

Release 18.3R1 Changes in Behavior and Syntax

There are no changes in behavior and syntax for NFX Series in Junos OS Release 18.3R1.

SEE ALSO

- New and Changed Features | 191
- Known Behavior | 193
- Known Issues | 194
- Resolved Issues | 195
- Documentation Updates | 196
- Migration, Upgrade, and Downgrade Instructions | 197
- Product Compatibility | 199

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 for the NFX Series.

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

NFX150 Series Devices

- The file transfer rate from an external media over the network to an NFX150 device is around 40– 50 Mbps. PR1290263
- On NFX150 devices running Junos OS Release 18.1R1, Transcend does not support Linux based SSD firmware upgrade mechanism in field for its SSD. Hence, field upgrade of Transcend SSD firmware cannot be provided for NFX150 devices. PR1347562

SEE ALSO

- New and Changed Features | 191
- Changes in Behavior and Syntax | 192
- Known Issues | 194
- Resolved Issues | 195
- Documentation Updates | 196
- Migration, Upgrade, and Downgrade Instructions | 197
- Product Compatibility | 199

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 for the NFX Series devices.

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

Known Issues: 18.3R2

- On NFX 150 devices,connectivity fault management (CFM) is not supported on circuit cross-connect (CCC) interfaces. PR1311588
- Under some circumstances, FPC0 ukern of NFX150 may crash and restart. The FPC recovers automatically and it does not crash again after the recovery. There is no known workaround. PR1347629
- On NFX 150 devices, dev key revocation is not supported by BIOS. Dev key revocation is to prevent customers from installing Dev signed image by mistake on their setup. PR1344738
- On NFX150 devices, syslog messages do not display xauth client authentication information such as assigned IP address and DNS. PR1305078
- On NFX150 devices, FTP displays an error message, **ftpd[14105]:bl_init: connect failed for `/var/run/blacklistd.sock'(No such file or directory)**. PR1315605
- On NFX150 devices, error messages are seen while rebooting the FPC0 interface. PR1326487
- During BIOS upgrade process, it does not display the existing BIOS version or the new BIOS version to which it is being upgraded. Similarly, it does not display the BIOS version when a lower version of BIOS is getting upgraded to a higher version of BIOS. PR1342573

SEE ALSO

- New and Changed Features | **191**
- Changes in Behavior and Syntax | **192**
- Known Behavior | **193**

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in the Junos OS Release 18.3R2 for NFX Series devices.

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

Resolved Issues: 18.3R2

- When ALG is enabled, some ALG traffic might fail to match the interface filter after packet reinjection. The reason is that **mbuf copy** command does not copy the filter information to new mbuf. As a workaround, disable the ALG. PR1339912
- Class super-user does not have group permission to create folder under /var/third-party. This allows non-root users of super-user class to create folders and files in /var/third-party. As a workaround, manually set the permission of /var/third-party through chmod 1777 /var/third-party. PR1352561
- On NFX250-S1E devices, the jsxe0 interface might not get IP address from the DHCP server. PR1354596
- NFX250 devices running Junos OS Release 18.1R1 do not provide product details to super-user when queried using netconf. As a workaround, log in as a root user. PR1356896
- The init.xml file is corrupted when there is an unexpected power outage, and the VNF fails to instantiate. When the init.xml file is corrupt or missing, the NFX device will recreate the init.xml file and instantiate the VNF by using the recreated init.xml file. As a workaround, check if the vnf instantiates after multiple power outages. PR1373997
- **NFX3/ACX5448:LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** message is displayed when you commit on configuration mode. As a workaround, run the

command, **set system syslog user * match "!(LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified)"** and commit. PR1376665

- The dialer-options route configuration in LTE interface does not work in user-defined routing instance. It works only in root routing instance. PR1389907

Resolved Issues: 18.3R1

There are no fixed issues in Junos OS Release 18.3R1 for the NFX Series.

SEE ALSO

- New and Changed Features | **191**
- Changes in Behavior and Syntax | **192**
- Known Behavior | **193**
- Known Issues | **194**
- Documentation Updates | **196**
- Migration, Upgrade, and Downgrade Instructions | **197**
- Product Compatibility | **199**

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 documentation for NFX Series.

SEE ALSO

- New and Changed Features | **191**
- Changes in Behavior and Syntax | **192**
- Known Behavior | **193**
- Known Issues | **194**
- Resolved Issues | **195**
- Migration, Upgrade, and Downgrade Instructions | **197**
- Product Compatibility | **199**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the NFX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

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.

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

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

Basic Procedure for Upgrading to Junos OS Release 18.3

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: 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 device, 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 device. For more information, see the Installation and Upgrade Guide..

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

To download and install Junos OS Release 18.3R1:

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

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

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

3. Select the **Software** tab.

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

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 device or to your internal software distribution site.

10. Install the new package on the device.

SEE ALSO

Product Compatibility

IN THIS SECTION

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 NFX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: https://apps.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.

Software Version Compatibility

This section lists the vSRX and Cloud CPE Solution software releases that are compatible with the Junos OS releases on the NFX Series devices.

NOTE: Starting in Junos OS Release 18.1R1, NFX150 and NFX250 devices support the same version of platform software and vSRX. For example, see Table 1 on page 200.

NFX250 Software Version Compatibility

This section lists the vSRX and CloudCPE Solution software releases that are compatible with the Junos OS releases on the NFX250 devices:

Table 1: Software Compatibility Details with vSRX and Cloud CPE Solution

NFX250 Junos OS Release	vSRX	Cloud CPE Solution
15.1X53-D40.3	15.1X49-D40.6	Cloud CPE Solution 2.0
15.1X53-D41.6	15.1X49-D40.6	Cloud CPE Solution 2.1
15.1X53-D102.2	15.1X49-D61	Cloud CPE Solution 3.0
15.1X53-D47.4	15.1X49-D100.6	Cloud CPE Solution 3.0.1
15.1X53-D490	15.1X49-D143	Cloud CPE Solution 4.0
15.1X53-D495	15.1X49-D160	Cloud CPE Solution 4.1
15.1X53-D496	15.1X49-D170	Cloud CPE Solution 4.1
15.1X53-D45.3	15.1X49-D61	Not applicable
17.2R1	15.1X49-D78.3	Not applicable
17.3R1	15.1X49-D78.3	Not applicable
17.4R1	15.1X49-D78.3	Not applicable
15.1X53-D471	15.1X49-D143	Not applicable
18.1R1	18.1R1	Not applicable
18.1R2	18.1R2	Not applicable
18.1R3	18.1R3	Not applicable
18.2R1	18.2R1	Not applicable

Table 1: Software Compatibility Details with vSRX and Cloud CPE Solution *(continued)*

NFX250 Junos OS Release	vSRX	Cloud CPE Solution
18.3R1	18.3R1	Not applicable
18.3R2	18.3R2	Not applicable

SEE ALSO

New and Changed Features | **191**

Changes in Behavior and Syntax | **192**

Known Behavior | **193**

Known Issues | **194**

Resolved Issues | **195**

Documentation Updates | **196**

Migration, Upgrade, and Downgrade Instructions | **197**

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the PTX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

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

Release 18.3R2 New and Changed Features

MPLS

- **Control transport address used for targeted-LDP session (PTX Series)**—Currently, only the router-ID or interface address is used as the LDP transport address. Starting in Junos OS Release 18.3R2, you can configure any other IP address as the transport address of targeted LDP sessions, session-groups, and interfaces. This new configuration is applicable only for configured LDP neighbors that have Layer 2 circuit, MPLS, and VPLS adjacencies.

  This feature is beneficial when you have multiple loopback interface addresses, and different IGPs associated with LDP interfaces, and you can control the session established between targeted LDP neighbors with the configured transport address.

  [See Control Transport Address Used for Targeted-LDP Session.]

Release 18.3R1 New and Changed Features

Hardware

- **PTX10K-LC1104 DWDM line card on PTX10016 routers**—Starting in Junos OS Release 18.3R1, the PTX10016 Packet Transport Router supports the PTX10K-LC1104 DWDM line card. The PTX10K-LC1104 DWDM line card provides packet forwarding speed up to 1.2 Tbps for cloud providers, service providers, and enterprises that need coherent dense wavelength-division multiplexing (DWDM) with MACsec security features. The six-port line card, with built-in optics, supports flexible rate modulation at 100-Gbps, 150-Gbps, and 200-Gbps speeds. The PTX10016 routers support a maximum of four PTX10K-LC1104 line cards.

  [See PTX10K-LC1104 Line Card.

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for password change policy enhancement (PTX Series)**—Starting in Junos OS Release 18.3R1, the Junos password change policy for local user accounts is enhanced to comply with certain additional password policies. As part of the policy improvement, you can configure the following:
  - **minimum-character-changes**—The number of characters by which the new password should be different from the existing password.
  - **minimum-reuse**—The number of older passwords, which should not match the new password.

  See password

Class of Service

- **Support for classification override (PTX Series)**—Starting with Junos OS 18.3R1, PTX Series devices support overriding the input classification of traffic streams, whether the input classification is the default forwarding class or some other forwarding class set by an input filter, assigning the streams to a different forwarding class by defining a policy class and referencing this class when configuring a routing policy.

To override the input classification, configure the **classification-override forwarding-class *class-name*** statement at the **[edit class-of-service forwarding-policy class *class-name*]** hierarchy level.

[See Overriding the Input Classification.]

- **RED drop and tail drop differentiation (PTX Series)**—Starting in Junos OS Release 18.2X75-D10 and Junos OS Release 18.3R1, PTX Series devices display traffic drops through tail drop counters unless there is an explicit RED drop profile configured. With a RED drop profile configured, traffic drops are reported under a RED drop counter. This configuration helps to differentiate between RED drop and tail drop scenarios.

  [See show interfaces queue.]

Interfaces and Chassis

- **Support for PTX10K-LC1104 coherent line card (PTX10008 and PTX10016)**—Starting with Junos OS Release 18.3R1, PTX10008 and PTX10016 routers support the PTX10K-LC1104 coherent DWDM line card. The following are some of the software features supported by this line card:
  - Compliance with ITU G.709 and ITU G.798
  - Performance-monitoring features such as alarms, threshold-crossing alarms, OTU/ODU error seconds, and FEC and bit error rate (BER) statistics
  - SNMP management of the MIC based on RFC 3591
  - IEEE 802.1ag OAM
  - IEEE 802.3ah OAM
  - IFINFO/IFMON
  - IEEE 802.3ad link aggregation
  - Flexible Ethernet services encapsulation
  - Flexible VLAN tagging
  - Source address MAC accounting per logical interface
  - Source address MAC filter per port
  - Source address MAC filter per logical interface
  - Destination address MAC filter per port
  - Up to 8000 logical interfaces shared across all ports on a single Packet Forwarding Engine

  See Understanding the PTX10K-LC1104 Line Card.

- **Support for BGP Monitoring Protocol on a nondefault management instance (PTX Series)**—Starting in Junos OS Release 18.3R1, the BGP Monitoring Protocol (BMP) can send monitoring packets to BMP monitoring stations that are reachable through a VRF table. This feature can be used with the **management-instance** configuration statement to create the routing instance mgmt-junos for BMP to

move through. Previously, BMP could send monitoring packets to a BMP monitoring station that could be looked up only using the default (inet.0 or inet6.0) routing table.

[See Configuring BGP Monitoring Protocol to Run Over a Different Routing Instance.]

Junos Telemetry Interface

- **Support for the Junos Telemetry Interface (ACX6360, MX Series, and PTX Series)**—Starting with Junos OS Release 18.3R1, Junos Telemetry Interface support is available for the ACX6360 Universal Metro Router and MX Series and PTX Series routers with a CFP2-DCO optics module that provides a high-density, long-haul optical transport network (OTN) transport solution with MAC capability.

  You can provision sensors to export telemetry data to an outside collector.

  The following native (UDP) and gRPC sensors can be provisioned for 100-Gigabit Ethernet interfaces and OTN interfaces:

  - **/junos/system/linecard/optical**
  - **/junos/system/linecard/otn**

  To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

  [See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **Expanded physical interface queue and traffic statistics sensors for Junos Telemetry Interface (JTI) (PTX, MX, EX, QFX, ACX)**—Starting with Junos OS Release 18.3R1, additional resource paths are added to stream physical (IFD) statistics.

  Prior to Junos OS Release 18.3R1, both traffic and queue statistics for physical interfaces (IFD) are sent out together using the resource path **/interfaces** for gRPC streaming (which is internally used to create **/junos/system/linecard/interface/**) or **/junos/system/linecard/interface/** for UDP (native) sensors.

  Now, traffic and queue statistics can be delivered separately. Doing so can reduce the reap time for non-queue data for platforms supporting Virtual Output Queues (VOQ).

  The following UDP resource paths can be configured:

  - **/junos/system/linecard/interface/** is the existing resource path (no change). Traffic and queue statistics are sent together.
  - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
  - **/junos/system/linecard/interface/queue/** exports queue statistics.

  The gRPC resource path **/interfaces** now has the following behavior:

  - In releases prior to Junos OS 18.3R1, it delivers all IFD traffic and queue statistics. In Junos OS 18.3R1 and higher, it delivers statistics in two sensors:

- **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
- **/junos/system/linecard/interface/queue/** exports queue statistics.

To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

[See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

- **Expanded ON_CHANGE support for LLDP telemetry data through Junos Telemetry Interface (JTI) (MX Series and PTX Series)**—Starting with Junos OS Release 18.3R1, OpenConfig support through remote procedure calls (gRPC) and JTI is expanded to support additional ON-CHANGE support for LLDP telemetry sensors. Periodical streaming of OpenConfig operational states and counters collects information at regular intervals. ON_CHANGE support streams operational states as events (only when there is a change), and is preferred over periodic streaming for time-sensitive missions.

  When you create a subscription using a top-level container as the resource path (for example, **/lldp**), leaves under the resource path **/lldp** with ON_CHANGE support are automatically streamed based on events. Other leaves will not be streamed.

  These paths, previously supporting periodical streaming only, now also support ON_CHANGE streaming:

  - /lldp/state/enabled
  - /lldp/state/chassis-id
  - /lldp/state/chassis-id-type
  - /lldp/state/system-name
  - /lldp/state/system-description
  - /lldp/state/hello-timer
  - /lldp/interfaces/interface/state/name
  - /lldp/interfaces/interface/state/enabled
  - /lldp/interfaces/interface/neighbors/neighbor/state/chassis-id
  - /lldp/interfaces/interface/neighbors/neighbor/state/chassis-id-type
  - /lldp/interfaces/interface/neighbors/neighbor/state/port-id
  - /lldp/interfaces/interface/neighbors/neighbor/state/port-id-type
  - /lldp/interfaces/interface/neighbors/neighbor/state/port-description

- /lldp/interfaces/interface/neighbors/neighbor/state/system-name
- /lldp/interfaces/interface/neighbors/neighbor/state/system-description
- /lldp/interfaces/interface/neighbors/neighbor/state/management-address
- /lldp/interfaces/interface/neighbors/neighbor/state/management-address-type
- /lldp/interfaces/interface/neighbors/neighbor/capabilities

These resource paths from the above list do not change with an event, but will be streamed on creation and deletion:

- /lldp/interfaces/interface/neighbors/neighbor/state/chassis-id
- /lldp/interfaces/interface/neighbors/neighbor/state/chassis-id-type
- /lldp/interfaces/interface/neighbors/neighbor/state/system-name

Before events are streamed, there is an initial stream of states to the collector, followed by an END_OF_INITIAL_SYNC notice. This notice signals the start of event streaming.

To provision the sensor to export data through gRPC streaming, use the **telemetry Subscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

[See Configuring a Junos Telemetry Interface Sensor (CLI Procedure) and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **BGP and statically configured SRTE traffic statistics sensor support for Junos Telemetry Interface (JTI) (MX Series and PTX Series)**—Starting with Junos OS Release 18.3R1, you can export traffic statistics for both ingress IP traffic and transit MPLS traffic that take segment routing traffic engineering (SRTE) paths. This feature provides support for BGP [draft-SRTE] and statically configured SRTE policies at ingress routers.

  JTI statistics can be exported using either gRPC streaming or UDP native sensors. The following resource paths are supported.

  For UDP native sensors:

  - **/junos/services/segment-routing/traffic-engineering/ingress/usage/**
  - **/junos/services/segment-routing/traffic-engineering/transit/usage/**

  For gRPC streaming:

  - **/mpls/signaling-protocols/segment-routing/**

  For exporting statistics by using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

  To provision the sensor to export data through gRPC streaming, use the **telemetry Subscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

For both export methods, you also must specify that these statistics be collected. To do this, configure collection at the **[edit protocols source-packet-routing telemetry statistics]** hierarchy level.

[See sensor, sensor, source-packet-routing and Guidelines for gRPC Sensors (Junos Telemetry Interface).

- **Bundling OpenConfig and Network Agent packages into Junos OS (MX Series, EX Series, PTX Series, ACX Series, QFX Series)**——OpenConfig and Network Agent packages are now bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

  OpenConfig can be found as a default package named **junos-openconfig**, and Network Agent content exists in the Junos OS as a daemon through the **na-telemetry** package. You can also install the OpenConfig package as an add-on package on top of the default package if you want to upgrade OpenConfig without upgrading Junos OS.

  [See Installing the OpenConfig Package, and Installing the Network Agent Package (Junos Telemetry Interface).]

MPLS

- **Support for next-hop-based dynamic UDP tunnels (PTX Series)**—Starting in Junos OS Release 18.3R1, next-hop-based MPLS-over-UDP dynamic tunnels are supported on the PTX Series routers. For every dynamic tunnel configured on a PTX router, a tunnel composite next hop, an indirect next hop, and a forwarding next hop is created to resolve the tunnel destination route. You can also use policy control to resolve the dynamic tunnel over select prefixes.

  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.

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

Network Management and Monitoring

- **Support for customized MIBs for sending custom traps based on syslog events (PTX Series)**—Starting in Junos OS Release 18.3R1, there is a process whereby customers can define their own MIBs for trap notifications. The customized MIB maps a particular error message with a custom OID rather than a generic one. Juniper Networks provides two new MIB roots reserved for customer MIBs, one for the custom MIB modules and the other for the trap notifications. For this process, you must convert the MIB to YANG format, and a tool is available for that.

  [See Customized SNMP MIBs for Syslog Traps.]

Routing Policy and Firewall Filters

- **Support added for logical and physical interface policers (PTX Series routers with third-generation FPCs)**—Starting with Junos OS Release 18.3R1, you can police both IPv4 and IPv6 traffic on the same link, at the physical and or logical interface level, by using an aggregate policer for both traffic types.

  As such, you can define a policer for a physical interface and then reference it in different firewall filters that can then be applied to a physical interface. For example, you can configure a single aggregate policer

for a physical interface, and then apply that policer to different logical interfaces and traffic families (IPv4 and IPv6) configured on the interface.

[See: Logical Interface (Aggregate) Policer Overview .]

- **Support to configure IPv6 packet flow labels as the load-balancing hash key (PTX Series routers)**—Starting in Junos OS Release 18.3R1, you can configure IPv6 packet flow labels for hash calculations on PTX Series routers. In releases before Junos OS Release 18.3R1, the **ipv6-flow-label** field in the IPv6 header is not used as the hashing key for load balancing. With this feature support, if you want the load balancing to be based on the flow label of the IPv6 header, include the **ipv6-flow-label** configuration statement at the **[edit forwarding-options hash-key family inet6 layer-3]** hierarchy level.

  [See ipv6-flow-label.]

Routing Protocols

- **Junos OS, OpenConfig, and Network Agent packages are delivered in a single TAR file (PTX Series)**—Starting in Junos OS Release 18.3R1, the Junos OS image includes the OpenConfig package and Network Agent; therefore, you do not need to install OpenConfig or Network Agent separately on your device.

  [See Installing the OpenConfig Package and Installing the Agent Network Package.]

Software Defined Networking (SDN)

- **Support for PCE-initiated point-to-multipoint LSPs (PTX Series)**—Starting in Junos OS Release 18.3R1, the Path Computation Element Protocol (PCEP) functionality is extended to allow a stateful PCE to initiate, provision, and modify point-to-multipoint traffic engineering LSPs through a PCC.

  Currently, Junos OS supports only point-to-point PCE-initiated LSPs. With the introduction of point-to-multipoint PCE-initiated LSPs, a PCE can initiate and provision a point-to-multipoint LSP dynamically without the need for local LSP configuration on the PCC. The PCE can also control the timing and sequence of the point-to-multipoint path computations within and across PCEP sessions, thereby creating a dynamic network that is centrally controlled and deployed.

  [See Understanding Path Computation Element Protocol for MPLS RSVP-TE with Support for PCE-Initiated Point-to-Multipoint LSPs.]

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (PTX Series)**—Starting in Junos OS Release 18.3R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database. The ephemeral database provides a fast programmatic interface that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. The active configuration of a device is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database. Ephemeral configuration data is volatile and is deleted when the device is rebooted.

  [See Understanding the Ephemeral Configuration Database.]

SEE ALSO

Changes in Behavior and Syntax | 211
Known Behavior | 214
Known Issues | 217
Resolved Issues | 220
Documentation Updates | 225

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for the PTX Series.

Interfaces and Chassis

- **Change in behavior for new port speed configuration on the PTX10K-LC1101 line card**—In Junos OS Release 18.3R1 and later, if you modify the port speed configuration on the PTX10k-LC1101 line card, then the new port speed configuration does not cause an FPC to reboot automatically, but it triggers an **FPC need bounce** alarm. To ensure that the new port speed configuration takes effect, you must manually reboot the FPC. The alarm is cleared when you manually reboot the FPC or delete the new port speed configuration.

  This change in behavior is also observed in Junos OS Releases 17.2X75-D102, 17.2R3, 17.4R2, 18.1R3, 18.1R2, 18.2X75-D10, 18.2R1, and later.

- On a PTX10K-LC1104 line card on the PTX10008 routers, the optical interface, when configured in QPSK modulation format with FEC mode SDFEC, the encoding scheme should be Differential. The Non Differential encoding scheme is not supported for such a configuration and Link will not come up if configured.

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (PTX Series)**—In Junos OS Release 18.3R2, the **show lacp interfaces | display xml** command displays a new XML tag element **<lacp-hold-up-state>**. The **<lacp-hold-up-state>** displays the time interval an interface holds

before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces were in a single **<lacp-hold-up-information>** XML tag. Now, for each interface it is displayed in a separate **<lacp-hold-up-information>** XML tag.

Junos OS XML API and Scripting

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (PTX Series)**—Starting in Junos OS Release 18.3R1, Junos OS does not support configuring an MD5 or SHA-1 checksum hash to verify the integrity of local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) scripts or support using an MD5 or SHA-1 checksum hash with the **op url *url* key** option to verify the integrity of remote op scripts.

Network Management and Monitoring

- **Junos OS does not support management of YANG packages in configuration mode (PTX Series)**—Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages using the **run** command in configuration mode is not supported.
- **Change in Error Severity (PTX10016)**—Starting in Junos OS Release 18.3R2, on PTX10016 routers, the severity of the FPC error, shown in the syslog as **PE Chip::FATAL ERROR!! from PE2[2]: RT: Clear Fatal if it is detected LLMEM Error MEM:llmem, MEMTYPE: 1**, is changed from fatal to non-fatal (or minor). In case of this error, only a message is displayed for information purpose. To view the error details, you can use the show commands **show chassis fpc errors** and **show chassis errors active**.

  [See show chassis fpc errors]
- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (PTX Series)**—Starting in Junos OS Release 18.3R2, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

Openconfig

- **The native Junos OS OpenConfig package does not display in operational mode (PTX Series)**—Starting in Junos OS Release 18.3R1, the **show system yang package** command in operational mode does not display the native Junos OS OpenConfig package. This is because the Junos OS image includes the OpenConfig package.

Routing Policy and Firewall Filters

- **Error caused by firewall filters with syslog and accept action (PTX1000 or PTX series routers with type 3 FPCs)**—In this release of Junos OS, under rare circumstances, the host interface may stop sending packets and the connections to and from the peer might fail if an outbound firewall filter is configured with an action of **syslog** and **accept**. This condition applies to IPv4 and IPv6 traffic families. Juniper recommends that you do not use the **syslog** and **accept** action in the output filter for these systems.

  An example configuration is provided (shows IPv4).

  ```
  set interfaces interface name unit unit family inet filter output name
  set firewall family inet filter name term 1 then syslog
  set firewall family inet filter name term 1 then accept
  ```

  [For more information, see PR 1354580.]

Software Installation and Upgrade

- **ssh-keygen output is tagged in XML (PTX1000)**—Starting in Junos OS Release 18.3R1, the output of the ssh-keygen utility that is invoked when generating the ssh keys, is now in XML format, and is wrapped in **<output>** tags. You can see this change in the console output at the time a device boots up with a new image.

  [See Junos OS Installation Package Names.]

Subscriber Management and Services

- **DHCPv6 lease renewal for separate identity association (IA) renew requests (PTX Series)**—Starting in Junos OS Release 18.3R1, the jdhcpd process handles the second renew request differently in the situation where the DHCPv6 client CPE device does both of the following:
  - Initiates negotiation for both the IA_NA and IA_PD address types in a single solicit message.
  - Sends separate lease renew requests for the IA_NA and the IA_PD and the renew requests are received back-to-back.

The new behavior is as follows:

1. When the reply is received for the first renew request, if a renew request is pending for the second address type, the client stays in the renewing state, the lease is extended for the first IA, and the client entry is updated.
2. When the reply is received for the second renew request, the lease is extended for the second IA and the client entry is updated again.

In earlier releases:

1. The client transitions to the bound state instead of staying in the renewing state. The lease is extended for the first IA and the client entry is updated.
2. When the reply is received for the second renew request, the lease is not renewed for the second address type and the reply is forwarded to the client. Consequently, when that lease ages out, the binding for that address type is cleared, the access route is removed, and subsequent traffic is dropped for that address or address prefix.

[See Using DHCPv6 IA_NA with DHCPv6 Prefix Delegation Overview.]

SEE ALSO

New and Changed Features | 202

Known Behavior | 214

Known Issues | 217

Resolved Issues | 220

Documentation Updates | 225

Migration, Upgrade, and Downgrade Instructions | 225

Product Compatibility | 230

Known Behavior

IN THIS SECTION

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 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.

General Routing

- In the specific case of semi-graceful RCB reboot initiated by the internal shell command **vhclient init 0**, GRES takes longer to complete i.r 3 minutes as opposed to 21 seconds. The regular CLI command **request vmhost reboot** (graceful) and a jack-out-jack-in of the Routing Engine (ungraceful) do not exhibit this delay. PR1312065
- Due to ZH ASIC limitation, MAC statistics under **show interface**, in Routing Engine, might not reflect "Mac error Counters" properly if ingress packet size is greater than default mtu (1518) or user configured mtu size (set interface mtu size <interface-name>mtu <288...9600>. PR1345779
- 100G DAC connected between QFX5200 and PTX10002-60C/QFX10002-60C will not link up. This is because BCM based devices have link-training enabled and Provider Edge-based devices do not have link-training enabled for 100G DAC/CR4. PR1356834
- Logical systems feature is not supported. PR1361016
- On a PTX1000 router, after the system is rebooted by issuing the **request vmhost reboot** command, the netproxy service might fail to start. PR1365664
- Frequent speed changes on interface ports might cause the relevant port IFD not being created. PR1367946

Interfaces and Chassis

- **Incorrect MAC statistics**—On the PTX10001 routers, the **show interfaces *interface-name* media detail** command might not display the MAC error counters correctly for oversized frames, that is, the ingress packets with size greater than the default maximum transmission unit (MTU) size (1518) or the MTU size configured by a user (by using the command **set interface interface-name mtu *bytes***).
- **The request support information command executes additional show commands**—On PTX Series routers, the **request support information** command executes the following show commands in addition to the existing **show** commands:
  - **show chassis fabric summary**
  - **show chassis fabric fpcs**
  - **show chassis fabric sibs**
  - **show chassis fabric topology**
  - **show chassis fabric reachability**

- **show chassis fpc**
- **show chassis power**
- **show pfe statistics traffic**

- On a PTX10K-LC1104 line card on the PTX10008 routers, the optical interface, when configured in QPSK modulation format with FEC mode SDFEC, the encoding scheme should be Differential. The link does not come up when configured with a non Differential encoding scheme as it is not supported.

Routing Policy and Firewall Filters

- **Error caused by firewall filters with syslog action (PTX1000 or PTX series routers with type 3 FPCs)**—In this release of Junos OS, under rare circumstances, the host interface may stop sending packets and the connections to and from the peer might fail if an outbound firewall filter is configured with an action of **syslog**. This condition applies to IPv4 and IPv6 traffic families. Juniper recommends that you do not use the **syslog** action in the output filter for these systems.

  An example configuration is provided (shows IPv4).

  ```
  set interfaces interface name unit unit family inet filter output name
  set firewall family inet filter name term 1 then syslog
  set firewall family inet filter name term 1 then accept
  ```

  [For more information, see PR 1354580.]

User Interface and Configuration

- **Auto-complete caution for QFX10002-60c and PTX10002-60c personalities**—Starting in Junos OS Release 18.3R2, for QFX10002-60c and PTX10002-60c personalities, do not use auto-complete to display the list of arguments for the **request system software delete** command. You must look for the package name using the **show system software** command and then explicitly type the software package name in the **request system software delete** command.

  [See request system software delete].

SEE ALSO

New and Changed Features | **202**

Changes in Behavior and Syntax | **211**

Known Issues | **217**

Resolved Issues | **220**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 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.

Interfaces and Chassis

- Upgrading Junos OS Release 14.2R5 and later maintenance releases and Junos OS Release 16.1 and later mainline releases with CFM configuration might cause the cfmd process to crash after upgrade. This is because of the old version of **/var/db/cfm.db**. PR1281073

General Routing

- Control packets might get dropped when the Packet Forwarding Engine experiences heavy congestion. PR1163759
- In a rare race condition, multiple interrupts are not handled properly on MX platform with MPC7E/MPC8E/MPC9E and PTX platform with FPC3-PTX-U2/FPC3-PTX-U3, which could lead to a core files. As a workaround, the interrupt code is optimized to avoid the unnecessary call to prevent the issue. PR1208536
- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09**

10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition errorPR1268678

- On the third-generation PTX Series routers FPCs (PTX3000/5000 FPC3, PTX1000) if the **protocols mpls no-propagate-ttl** command is configured, the MPLS TTL field can be reset to 255 in the packets where a label swap operation is performed. PR1287473
- On a PTX Series PIC with the CFP2-DCO-T-WDM transceiver installed, after repeated configuration rollback, the link sometimes takes a long time to come up. PR1301462
- Repeated log message **%PFE-3 fpcX expr_nh_index_tree_ifl_get** and **expr_nh_index_tree_ipaddr_get** are observed when sampling packets are discarded with log(or syslog) statements under the firewall filter. PR1304022
- On a PTX Series router with a third-generation FPC, the error message is displayed when the FPC goes online or offline. PR1322491
- On 30-Port MacSec Linecard (LC1101-M - 30C / 30Q / 96X) of Vale-PTX chassis, under certain circumstances, while **exclude-protocol lacp** configuration under Hierarchy Level **[edit security macsec connectivity-association connectivity-association-name]** is deleted or deactivated, the LACP Protocol's "Mux State" shown under the output of CLI command **show lacp interface**, might remain as "attached" or "detached" and would not transition to "distributing" state. PR1331412
- The output of the CLI command **show class-of-service fabric statistics** now includes traffic that was dropped because of internal errors in the drop counts. PR1338647
- PTX3000 Series routers report CCL (Chip to Chip Link) CRC errors while FPC3-SFF-PTX-1X is offlined through CLI command or by pressing the offline button. The syslog error is generated by an FPC just before it goes offline, so there is no detectable traffic loss. **Apr 2 08:43:00 fpc4 CMSNGFM: cmsngfpc_fm_send_spry_ctrl_ack: ev_id:11 fm_st:ALL fm_type:FPC_OFF fm_op:DEL Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 0 detected link error for S00F0_0(11,0,11)->FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: Logging statistics for FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: SOT:0x0000037649c2c43e Apr 2 08:43:00 fpc2 CCL: FrameCnt:0x00000000000419dc Apr 2 08:43:00 fpc2 CCL: LastCRCErrCnt:0x00000003 Apr 2 08:43:00 fpc2 CCL: AggrCRCErrCnt:0x0000000000000003 Apr 2 08:43:00 fpc2 CCL: AggrBERCnt:0x0000000000000001 Apr 2 08:43:00 fpc2 CCL: pe0-Avg-28nm-link-10-18 CRC error history (last 5 polls): Apr 2 08:43:00 fpc2 CCL: 0x0 0x0 0x0 0x0 0x3 Apr 2 08:43:00 fpc2 CCL: FEC Uncorrectable FEC Correctable Apr 2 08:43:00 fpc2 CCL: 00000004, 00000000 Apr 2 08:43:00 fpc2 CCL: 00000000, 00000000 Apr 2 08:43:00 fpc2 BEGIN Rx serdes info for asic pe0-0 serdes 18 Apr 2 08:43:00 fpc2 Signal & port condition for serdes_num 18 Apr 2 08:43:00 fpc2 Rx Signal : Signal Not OK Apr 2 08:43:00 fpc2 Rx Electrical Idle : High Apr 2 08:43:00 fpc2 Rx Frequency Lock: Set Apr 2 08:43:00 fpc2 Rx Port : Ready Apr 2 08:43:00 fpc2 DFE TAPs : -- snip -- Apr 2 08:43:00 fpc2 CCL: FrameCnt:0x0000000000041a0d Apr 2 08:43:00 fpc2 CCL: LastCRCErrCnt:0x00000003 Apr 2 08:43:00**

fpc2 CCL: AggrCRCErrCnt:0x0000000000000003 Apr 2 08:43:00 fpc2 CCL: AggrBERCnt:0x0000000000000001 Apr 2 08:43:00 fpc2 CCL: pe0-Avg-28nm-link-14-22 CRC error history (last 5 polls): Apr 2 08:43:00 fpc2 CCL: 0x0 0x0 0x0 0x0 0x3 Apr 2 08:43:00 fpc2 CCL: FEC Uncorrectable FEC Correctable Apr 2 08:43:00 fpc2 CCL: 00000004, 00000000 Apr 2 08:43:00 fpc2 CCL: 00000000, 00000000 Apr 2 08:43:00 fpc2 BEGIN Rx serdes info for asic pe0-0 serdes 22 Apr 2 08:43:00 fpc2 Signal & port condition for serdes_num 22 Apr 2 08:43:00 fpc2 Rx Signal : Signal Not OK Apr 2 08:43:00 fpc2 Rx Electrical Idle : High Apr 2 08:43:00 fpc2 Rx Frequency Lock: Set Apr 2 08:43:00 fpc2 Rx Port : Ready Apr 2 08:43:00 fpc2 DFE TAPs : -- snip -- Apr 2 08:43:00 fpc2 CCL: Logging errors for FPC02FE0(0,00) Apr 2 08:43:00 fpc2 CCL: BER Err Apr 2 08:43:00 fpc2 CCL: Frame Lock Loss Apr 2 08:43:00 fpc2 CCL: Align Loss Apr 2 08:43:00 fpc2 CCL: Header Comparison Error Apr 2 08:43:00 fpc2 CCL: Header Preamble Error Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 0 detected link error for S00F1_0(14,0,14)->FPC02FE0(1,00) Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 1 detected link error for S00F0_0(11,0,11)->FPC02FE1(0,00) Apr 2 08:43:00 fpc2 CMSNGFM: cmsngfpc_platform_fm_periodic: PFE 1 detected link error for S00F1_0(14,0,14)->FPC02FE1(1,00) User@PTX3000> show chassis hardware detail Hardware inventory: FPC 0 REV 43 750-057064 ACPV7514 FPC3-SFF-PTX-1XCPU BUILTIN BUILTIN SMPC PMB FPC 2 REV 40 750-057064 ACPJ9145 FPC3-SFF-PTX-1XCPU BUILTIN BUILTIN SMPC PMB FPC 4 REV 43 750-057064 ACPR8506 FPC3-SFF-PTX-1XCPU BUILTIN BUILTIN SMPC PMB SIB 0 REV 10 750-057067 ACPJ8829 SIB3-SFF-PTX SIB 1 REV 10 750-057067 ACPJ8683 SIB3-SFF-PTX SIB 2 REV 10 750-057067 ACPJ8843 SIB3-SFF-PTX SIB 3 REV 10 750-057067 ACPJ8920 SIB3-SFF-PTX PR1348733

- With TIC offline or online, MPLS bidirectional traffic flow might stop working. PR1367920
- Unsuccessful connection attempts will not be logged on the backup SPMB. PR1369731
- On PTX Series platforms, if the aggregated Ethernet child interfaces are across different Packet Forwarding Engines and **nexthop-learning** is configured, the MAC filter statistics of the child interface might be abnormal. PR1370062
- On all PTX-Series or QFX10002/QFX10008/QFX10016 platforms with CoS deployed, all the physical member interfaces of aggregated Ethernet (AE) might drop the packets in lower priority queues when micro-bursts are received. These micro-burst are typically due to the speed differential between ingress interface (for example, 100G) and egress interface (for example, 10G). Typically it occurs when a large burst of high priority traffic and lower priority traffic arrive simultaneously. PR1385454
- The DHCP Relay functionality does not work on PTX10001-20C devices. DHCP relay functionality: The DHCP requests and the DHCP offers are snooped by the box, the snooping happens via firewall, firewall snoops all the DHCP packets ingressing the default route table and all the offers and requests are punted unto the host/control-plane. When a DHCP client sends the DHCP request, it gets intercepted by the filter block and punted up to the control plane. Upon receiving this packet, control-plane unicast (relay) this packet to DHCP server. DHCP server responds back with a DHCP Offer, which again gets intercepted by the firewall block and punted up. Upon receiving the DHCP offer, control plane broadcast this DHCP offer to the clients VLAN and eventually client receives the DHCP offer. PR1407476

- On PTX or QFX10002/QFX10008/QFX10016, a auto correctable non-fatal hardware error on PE chip (that is, ASIC on PTX1000, PTX10002, QFX10002, the third-generation FPC on PTX3000/PTX5000, and the Line card on PTX10008/PTX10016/QFX10008/QFX10016) is reported as 'FATAL' error and hence the related Packet Forwarding Engine (PFE) will get disabled. The code changes have been made to change the error category from 'FATAL' to 'INFO' to avoid the Packet Forwarding Engine to be disabled unexpectedly. PR1408012
- The **rx_power** value streamed to the telemetry server is the raw value (mW) returned directly from the transceiver driver. The Junos CLI value has been transformed in the transportd daemon into different units: (Rx input total power(0.01dBm). PR1411023
- On PTX platform, hostname does not update at FPC shell after host name change unless FPC reboots. PR1412318

Routing Protocols

- When the loopback interface is configured in a logical-system and Routing Engine-based micro BFD is configured to use the loopback address as source address, BFD packets go out with source address belonging to outgoing interface rather than the loopback address. Due to this issue, the micro BFD session might not be able to come up. PR1370463

SEE ALSO

New and Changed Features | **202**

Changes in Behavior and Syntax | **211**

Known Behavior | **214**

Resolved Issues | **220**

Documentation Updates | **225**

Migration, Upgrade, and Downgrade Instructions | **225**

Product Compatibility | **230**

Resolved Issues

IN THIS SECTION

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

Resolved Issues: 18.3R2

General Routing

- Status LED on the chassis does not show UP on PTX10002-60C. PR1332991
- Member of IPv4 unilist next-hops might be stuck in "Replaced" state after interface flaps. PR1336201
- Multicast traffic packet drop of more than 50 percent is seen when having FPC1/FPC2 mix with FPC3. PR1339481
- Traffic drop might be seen after GRES if uRPF is configured. PR1354285
- The host interface might stop sending packets on PTX Series routers with FPC3 or PTX1000 when using outbound firewall filter with syslog option. PR1354580
- PTX10001 platform: FRR link-protection convergence time is not within limits. PR1355953
- Netproxy service client component fails start after issuing **request vmhost reboot**. PR1365664
- The 'Normal discards' Packet Forwarding Engine statistics traffic counter might increase at a higher rate when Inline-Jflow or sFlow is enabled. PR1368208
- The IPLC card might take a long time to come up. PR1368637
- Inline BFD might keep flapping when inline sampling is configured. PR1376509
- Traffic might be dropped on third-generation FPCs on PTX Series routers. PR1378392
- With the given configuration, continuous error messages are observed: **expr_sensor_detach_rt_cntr_from_sgid:2793: Invalid sensor group ID 0**.PR1379730
- FPC might crash on PTX Series routers or QFX10000 after lo0 filter change. PR1380917
- BFD sessions flap when restarting one FPC on PTX10000 routersPR1383703
- Packet Forwarding Engine-based local repair does not happen for IP routes pointing to unilist of composites with Indirect nexthops. PR1383965
- CPSM daemon memory leak is observed in VMHOST. PR1387903
- BFD flaps are seen on PTX Series routers or QFX10000 platforms with inline BFD. PR1389569
- Forwarding issue on mixed link-speed aggregated Ethernet interface after FPC reloads. PR1390417
- lcmd core and FPC restarted. PR1391443
- High jsd or na-grpcd CPU usage might be seen even when JET or JTI is not used. PR1398398
- CPU overuse might be observed on PTX/QFX10000 Series platform. PR1399369

- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. PR1399683
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- jlock hog by **tcp_timer_delack** caused ppmd connection to drop with Packet Forwarding Engine. PR1401507
- The log message **JAM HW data base open failed for ptx5kpic_3x400ge-cfp8** is seen during commit. PR1403071
- Incorrect mem stat message is seen in FPC logs of PTX Type 1 FPC. PR1404088
- PTX3000: FPCs are not able to come online for tens of minutes after a reboot of the chassis. PR1404611
- 100G SR4 Optics with part number 740-061405 should be displayed as "QSFP-100G-SR4-T2". PR1405399
- No chassis alarm is raised on PTX1000 when PEM is removed or power lost to PEM. PR1405430
- Layer2 VPN will flap repeatedly after link up between PE and CE under "asynchronous-notification" and "some types of MICs" conditions. PR1407345
- The L2circuit egress PE might drop the traffic in FAT+CW enabled L2circuit scenario when another FAT+CW enabled L2circuit PW flaps. PR1415614
- Traffic loss might be seen for duration of hold-time down timer when flapping an interface with hold-time down timer configured. PR1418425

Infrastructure

- The FPC might go down on some vmhost-based PTX platforms. PR1367477
- The error of **jlaunchd: disk-monitoring is thrashing, not restarted** might be seen. PR1380032

Interfaces and Chassis

- **PE Chip:pe0[0]: IPW: oversize_drop error** causes major error on FPC. PR1375030

MPLS

- MPLS LSP will keep down state due to routing loop detection after flapping link between P router and egress PE. PR1384929
- The rpd might crash when LDP route with indirect next-hop is deleted. PR1398876
- A single-hop bypass LSP might not be used for traffic when both transit chaining mode and sensor-based-stats are used. PR1401152

Platform and Infrastructure

- Some files are missing during log archiving. PR1405903

Routing Protocols

- The rpd process might crash after executing commit the configuration related to mapping-server-entry. PR1379558
- The rpd generates core files on backup Routing Engine during neighborship flap when using an authentication-key with size larger than 20 characters. PR1394082
- Syslog message is seen whenever prefix sid coincides with the node sid. PR1403729
- Memory leaks when labeled-isis transit routes created as chain composite nexthop. PR1404134

Resolved Issues: 18.3R1

General Routing

- On a PTX1000 router, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3 fails frequently when sampling is enabled. PR1296533
- On PTX10008 and PTX10016 routers, the suppress chassis alarm is displayed for switched-off PEMS. PR1311574
- On a PTX10000 router, when the PIC or port speed is changed in configuration, an alarm or warning is issued. PR1311875
- On a PTX10000 router, for 100-gigabit LR4 optics with part number 740-061409 change the **show chassis hardware** display to QSFP-100G-LR4-T2. PR1322082
- The GRE traffic is not decapsulated by the firewall filter. PR1325104
- On PTX10016 router, VXLAN is not supported and it has to be either removed or disabled in the **[edit switch-options]** hierarchy. PR1328502
- A next-hop programming issue occurs during link flapping on PTX Series routers. PR1333274
- The **Tc_count** counters in a filter configured with the **scale-optimized** statement do not increment. PR1334580
- On PTX10016 router, traffic stops flowing out of the aggregated Ethernet interface 70 after some FPC restarted a few times. PR1335118
- The FPC might get rebooted a few minutes after a configuration is loaded. PR1346467
- No DHCP service or configuration is running after the system is zeroized. PR1347730
- Sensors are not getting cleared up after a Routing Engine switchover is performed. PR1347779
- The MPLS traceroute for P2MP LSPs configured with link protection causes the FPC to crash. PR1348314
- On PTX Series routers, the **threshold** is not getting configured correctly when it is configured using **scope** and **category** options. PR1350841

- BFD sessions do not come up on PTX3000 router. PR1352112
- If the 15-port 100-Gigabit Ethernet PIC is used on PTX Series routers, the interface might be delayed by 60 seconds to come up from the down status. PR1357410
- The multicast replication traffic might be lost on an aggregated Ethernet bundle interface after one member link goes down. PR1359974
- Traffic continues to be forwarded through the member links of an aggregated Ethernet bundle interface even when the **Link-Layer-Down** flag is set. PR1365263
- On PTX IPLC (OPT3-SFF-PTX FPC), the first J-UKERN crash triggers multiple secondary J-UKERN crashes. PR1365791
- The **commit** or **commit check** command might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- Packet is dropped after the filter is deleted on the interface. PR1372957
- Layer 3 VPN traffic might be dropped because of the selector weight is set to 65535 after one core-facing interface goes down. PR1380783

MPLS

- MPLS LSP statistics are not shown in the output of the **show mpls lsp ingress statistics** command. PR1344039
- Some LSPs might be stuck on the upstream devices after interfaces flap occurs on downstream devices. PR1349157
- IPv6 routes are dead in the mpls.0 table and S=0 leads traffic loss in v6-indirect next-hop stitching. PR1355878
- LSP with auto-bandwidth enabled goes down during an HMC error condition. PR1374102

Platform and Infrastructure

- Running RSI through the console port might cause the system to crash and reboot. PR1349332
- Traffic might be silently dropped and the following error message for next hops is displayed: **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region**. PR1357707
- Unable to commit Junos OS configuration during the ZTP process and ZTP process stop is completed. PR1358919

Routing Protocols

- All OSPF neighbors go down after start the LDP session. PR1304504
- The primary path of an MPLS LSP might switch to other address. PR1316861
- Protocol churn might cause the rpd to crash. PR1341466
- An rpd core file might be generated while running streaming telemetry. PR1347431

SEE ALSO

New and Changed Features | 202
Changes in Behavior and Syntax | 211
Known Behavior | 214
Known Issues | 217
Documentation Updates | 225
Migration, Upgrade, and Downgrade Instructions | 225
Product Compatibility | 230

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 documentation for PTX Series.

SEE ALSO

New and Changed Features | 202
Changes in Behavior and Syntax | 211
Known Behavior | 214
Known Issues | 217
Resolved Issues | 220
Migration, Upgrade, and Downgrade Instructions | 225
Product Compatibility | 230

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Basic Procedure for Upgrading to Release 18.3 | 226
- Upgrade and Downgrade Support Policy for Junos OS Releases | 228
- Upgrading a Router with Redundant Routing Engines | 229

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 might take several hours, depending on the size and configuration of the network.

Basic Procedure for Upgrading to Release 18.3

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 Installation and Upgrade Guide.

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

To download and install Junos OS Release 18.3R2:

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

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

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

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

4. Select the **Software** tab.

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

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

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

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

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-18.3R2.9.tgz
```

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-18.3R2.9-limited.tgz
```

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

- ***/pathname***—For a software package that is installed from a local directory on the router.
- For software packages that are downloaded and installed from a remote location:
  - **ftp://*hostname/pathname***
  - **http://*hostname/pathname***

- **scp://*hostname*/*pathname***

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

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

Rebooting occurs only if the upgrade is successful.

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

NOTE: After you install a Junos OS Release 18.3 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

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 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from

Junos OS Release 17.1 to Release 17.3. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Upgrading a Router with Redundant Routing Engines

If the router has two Routing Engines, perform 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.

SEE ALSO

- New and Changed Features | 202
- Changes in Behavior and Syntax | 211
- Known Behavior | 214
- Known Issues | 217
- Resolved Issues | 220
- Documentation Updates | 225
- Product Compatibility | 230

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

To determine the features supported on PTX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at: https://apps.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 | 202
Changes in Behavior and Syntax | 211
Known Behavior | 214
Known Issues | 217
Resolved Issues | 220
Documentation Updates | 225
Migration, Upgrade, and Downgrade Instructions | 225

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the QFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

New and Changed Features

IN THIS SECTION

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

NOTE: The following QFX Series platforms are supported in Release 18.3R2: QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016.

Release 18.3R2 New and Changed Features

MPLS

- **Control transport address used for targeted-LDP session (QFX Series)**—Currently, only the router-ID or interface address is used as the LDP transport address. Starting in Junos OS Release 18.3R2, you can configure any other IP address as the transport address of targeted LDP sessions, session-groups, and interfaces. This new configuration is applicable only for configured LDP neighbors that have Layer 2 circuit, MPLS, and VPLS adjacencies.

  This feature is beneficial when you have multiple loopback interface addresses, and different IGPs associated with LDP interfaces, and you can control the session established between targeted LDP neighbors with the configured transport address.

  [See Control Transport Address Used for Targeted-LDP Session.]

Release 18.3R1-S3 New and Changed Features

- **Host route generation support for ARP and Neighbor Discovery Protocol (NDP) (QFX5100)**—Starting in Release 18.3R1-S3, Junos OS supports host route generation for devices connected to QFX5100 switches in a data center. When you enable this feature on an interface for IPv4 or IPv6, host routes are created in the routing table for each device present in ARP (IPv4) and NDP (IPv6). These host routes can be exported to routing protocols to be advertised to the network by matching the new policy qualifier **l2-learned-host-routing** statement.

  You can configure the **host-route-generation** statement under the **[edit interfaces *name* unit *name* family inet/inet6]** hierarchy, on each interface and for each address family.

  NOTE: Host route generation is disabled by default.

- **Proactive ARP detection (QFX5100)**—Starting with Junos OS Release 18.3R1-S3, you can check the reachability of connected devices (within an IP subnet range) on a specified interface.

  To enable proactive ARP detection, configure the **proactive_arp_detection** statement at the **[edit system arp]** hierarchy level. After you enable proactive ARP detection, an ARP request is sent over the interface, and the ARP reply received is updated in the ARP cache.

- **Layer 3 VXLAN gateway (QFX5120 switches)—**Starting in Junos OS Release 18.3R1-S3, you can deploy QFX5120 switches as a Layer 3 VXLAN gateway in EVPN-VXLAN overlay networks with the following IP fabric architectures:
  - A two-layer IP fabric that includes spine devices (Layer 3 VXLAN gateways) and leaf devices (Layer 2 VXLAN gateways). You can deploy QFX5120 switches as spine or leaf device in this fabric.
  - A one-layer IP fabric that includes leaf devices that function as both Layer 2 and Layer 3 VXLAN gateways. You can deploy QFX5120 switches as leaf nodes in this fabric.

  The QFX5120 switches also support EVPN pure type-5 routes.

  [See Understanding EVPN with VXLAN Data Encapsulation and Understanding EVPN Pure Type-5 Routes.]

Release 18.3R1-S2 New and Changed Features

EVPN

- **VXLAN support (QFX5120 switch)—**Starting in Junos OS Release 18.3R1-S2, the QFX5120 switch supports the following Virtual Extensible LAN (VXLAN) features:
  - Ethernet VPN-VXLAN (EVPN-VXLAN)—Layer 2 VXLAN gateway functionality.
    - Proxy Address Resolution Protocol (ARP) and ARP suppression, and Neighbor Discovery Protocol (NDP) and NDP suppression on non-IRB interfaces. [See EVPN Proxy ARP and ARP Suppression, and NDP and NDP Suppression.]
    - Internet Group Management Protocol (IGMP) snooping. [See Overview of IGMP Snooping in an EVPN-VXLAN Environment.]
    - Tunneling of Q-in-Q traffic. [See Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network].
    - MAC mobility. [See Overview of MAC Mobility.]
  - Physical Interface Module (PIM)-based Layer 2 VXLAN gateway. [See Examples: Manually Configuring VXLANs on QFX Series and EX4600 Switches.]

Release 18.3R1 New and Changed Features

Hardware

- **QFX5120-48Y switches**—Starting with Junos OS Release 18.3R1, the QFX5120-48Y switch is available as a fixed-configuration switch with the following built-in ports:
  - Forty-eight 25-Gigabit Ethernet ports that can operate at 1-Gbps, 10-Gbps, or 25-Gbps speed and support SFP, SFP+, or QSFP28 transceivers.

- Eight 100-Gigabit Ethernet ports that can operate at 40-Gbps or 100-Gbps speed and support QSFP+ or QSFP28 transceivers. When these ports operate at 40-Gbps speed, you can configure four 10-Gbps interfaces and connect breakout cables, increasing the total number of supported 10-Gbps ports to 80. When these ports operate at 100-Gbps speed, you can configure four 25-Gbps interfaces and connect breakout cables, increasing the total number of supported 25-Gbps ports to 80.

A total of four models are available: two featuring AC power supplies and front-to-back or back-to-front airflow and two featuring DC power supplies and front-to-back or back-to-front airflow.

CAUTION: QFX5120 switches require the use of Junos OS Release 18.3R1.11 which is available on the QFX5120 software download page.

[See QFX5120 Documentation.]

- **QFX5210-64C-DC switches**—Starting in Junos OS Release 18.3R1, Juniper Networks expands the QFX5210-64C line of switches to include DC power options. Like the existing AC models, the QFX5210-64C-DC is a 64-port, fixed-chassis switch designed for spine-and-leaf applications that need high-port density in next-generation IP fabric networks. All 64 ports in the 2 U, standalone switch default to 100 Gbps speeds but you can also configure the ports for 10 Gbps, 25 Gbps, 40 Gbps, and 50 Gbps speeds. The routing engine and control plane are driven by the 2.2 GHz quad-core Intel® Xeon® CPU with 16 GB of memory and an enterprise grade 100 GB solid-state drive (SSD) for storage. The QFX5210-64C-DC comes standard with redundant fans and redundant power supplies. The switch can be ordered with either ports-to-FRUs or FRUs-to-ports airflow.

  [See QFX5210 System Overview.]

- **Support for JNP-SFP-10G-BX40D and JNP-SFP-10G-BX40U transceivers (QFX5110)**— Starting in Junos OS Release 18.3R1, the QFX5100 switches support the JNP-SFP-10G-BX40D and the JNP-SFP-10G-BX40U transceivers. The JNP-SFP-10G-BX40D and JNP-SFP-10G-BX40U transceivers are for single SMF bidirectional applications. A JNP-SFP-10G-BX40D transceiver should always be connected to a JNP-SFP-10G-BX40U transceiver with a single SMF. The operating link distance is up to 40 km. With a single LC receptacle, the JNP-SFP-10G-BX40D transceiver transmits a 1330 nm wavelength signal and receives a 1270 nm signal, whereas JNP-SFP-10G-BX40U transceiver transmits a 1270 nm wavelength signal and receives a 1330 nm signal. [See the Hardware Compatibility Tool.]

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for password change policy enhancement (QFX Series)**—Starting in Junos OS Release 18.3R1, the Junos password change policy for local user accounts is enhanced to comply with certain additional password policies. As part of the policy improvement, you can configure the following:
  - **minimum-character-changes**—The number of characters by which the new password should be different from the existing password.
  - **minimum-reuse**—The number of older passwords, which should not match the new password.

See password

Class of Service (CoS)

- **Support for CoS on QFX5120 switches (QFX5120)**—Starting in Junos OS Release 18.3R1, the QFX5120 switch supports class of service (CoS) functionality. CoS is the assignment of traffic flows to different service levels. You can use CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to ensure quality of service (QoS) to particular applications served by specific traffic flows across the network.

  Compared to CoS functionality on QFX5100 and QFX5110 switches, QFX5120 switches provide significantly more buffer memory (32 MB), but do not support hierarchical scheduling or ETS. The QFX5120 also supports eight unicast and two multicast queues.

  [See CoS Support on QFX Series Switches, EX4600 Line of Switches, and QFabric Systems.]

EVPN

- **IPv4 inter-VLAN multicast forwarding modes for EVPN (QFX10000 switches)—**Starting with Junos OS Release 18.3R1, QFX10000 switches can forward IPv4 multicast traffic between VLANs in EVPN-VXLAN networks with these IP fabric architectures:

  - Two-layer IP fabric in which QFX10000 switches function as Layer 3 gateways, and QFX5100 or QFX5200 switches function as Layer 2 gateways.
  - One-layer IP fabric in which QFX10000 switches function as both Layer 2 and Layer 3 gateways.

    In both architectures, QFX10000 switches on which IRB interfaces are configured can route multicast traffic from one VLAN to another.

  [See Multicast Support in EVPN-VXLAN Overlay Networks.]

- **Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network (QFX10000 switches)**—Starting with Junos OS Release 18.3R1, QFX10000 switches that function as Layer 3 and Layer 2 virtual tunnel endpoints (VTEPs) can tunnel single-tagged and double-tagged Q-in-Q packets through an EVPN-VXLAN overlay network. In addition to tunneling Q-in-Q packets, the ingress and egress VTEPs can perform the following Q-in-Q actions:

  - Delete, or pop, an outer service provider VLAN (S-VLAN) tag from an incoming packet.
  - Add, or push, an outer S-VLAN tag onto an outgoing packet.
  - Map a configured range of customer VLAN (C-VLAN) IDs to an S-VLAN.

  > **NOTE:** The QFX10000 switches do not support the pop and push actions with a configured range of VLANs.

  [See Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network.]

- > **NOTE:** This feature is documented but not supported on QFX5110 switches in Junos OS Release 18.3R1.

  IPv6 data traffic support through an EVPN-VXLAN overlay network (QFX5110 switches)—Starting with Junos OS Release 18.3R1, QFX5110 switches that function as Layer 3 VXLAN gateways can route IPv6 data traffic through an EVPN-VXLAN overlay network. With this feature enabled, Layer 2 or 3 data packets from one IPv6 host to another IPv6 host are encapsulated with an IPv4 outer header and transported over the IPv4 underlay network. The Layer 3 VXLAN gateways in the EVPN-VXLAN overlay network learn the IPv6 routes through the exchange of EVPN type-2 and type-5 routes.

  [See Routing IPv6 Data Traffic through an EVPN-VXLAN Network With an IPv4 Underlay.]

- **Firewall filtering and policing on IRB Interfaces in EVPN-VXLAN (QFX10000 switches)**—Starting with Junos OS Release 18.3R1, you can configure a firewall filter on an IRB interface in an EVPN-VXLAN topology. The IRB interface acts as a Layer 3 routing interface to connect the VXLANs in one-layer or two-layer IP fabric topologies. Firewall filters can only be configured on the IRB interface after the VXLAN header is stripped by the VXLAN tunnel endpoint (VTEP). Only ingress filtering is supported.

  [See Firewall Filter Match Conditions and Actions for QFX10000 Switches.]

General Routing

- **Layer 3 unicast features (QFX5120)**—Starting with Junos OS Release 18.3R1, the following Layer 3 unicast features are supported:
  - Static routing, ping, and traceroute (IPv4, IPv6)
  - OSPFv2 (IPv4) and OSPFv3 (IPv6)
  - RIPv2
  - BGP (IPv4, IPv6), BGP 4-byte ASN support, and BGP multipath
  - MBGP (IPv4)
  - IS-IS (IPv4, IPv6)
  - BFD (for RIP, OSPF, IS-IS, BGP, PIM)
  - Unicast reverse path forwarding (RPF)
  - Filter based forwarding (FBF)
  - IP directed broadcast traffic forwarding
  - IPv4 over GRE
  - Virtual router redundancy protocol (VRRP)
  - VRRPv3 (IPv6)
  - Neighbor Discovery Protocol (IPv6)

- Path MTU discovery
- IPv6 class of service—Behavior aggregate (BA) classifiers, multifield (MF) classifiers and rewrite rules, traffic-class scheduling)
- IPv6 stateless address autoconfiguration
- Equal-cost multipath (ECMP)—32-way
- Virtual router (VRF-lite) IS-IS, RIP, OSPF, BGP

Interfaces and Chassis

- **Multichassis link aggregation groups, configuration synchronization, and configuration consistency check (MC-LAG) (QFX5120 switches)**—Starting with Junos OS Release 18.3R1, MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running Spanning Tree Protocol (STP).

  Configuration synchronization enables you to easily propagate, synchronize, and commit configurations from one MC-LAG peer to another. Log in to either peer to manage both, and use configuration groups to simplify the configuration process. You can create one configuration group each for the local peer and the remote peer, and a global configuration common to both peers. Create conditional groups to specify when peer configurations are synchronized.

  Use configuration consistency checks, which are enabled by default, to find configuration-parameter inconsistencies between multichassis link aggregation group (MC-LAG) peers.

  [See Multichassis Link Aggregation Features, Terms, and Best Practices].

- **Channelizing Interfaces on QFX5120-48Y Switches**—On the QFX5120-48Y switch, there are a total of 56 ports. Of these 56 ports, 8 ports (labeled 48 through 56) are uplink ports that support 100-Gigabit Ethernet interfaces (QSFP28 ports) and 40-Gigabit Ethernet interfaces (QSFP+ ports). The other 48 ports (labeled 0 through 47) are SFP+ ports that support 25-Gigabit Ethernet interfaces or 10-Gigabit Ethernet interfaces. The default speed for the SFP+ ports is 10 Gbps.

  Starting with Junos OS Release 18.3R1, you can channelize the 100-Gigabit Ethernet interfaces to four independent 25-Gigabit Ethernet interfaces. The default 100-Gigabit Ethernet interfaces can also be configured as 40-Gigabit Ethernet interfaces, and in this configuration can either operate as dedicated 40-Gigabit Ethernet interfaces, or can be channelized to four independent 10-Gigabit Ethernet interfaces using breakout cables on the QFX5120-48Y switch.

  > **NOTE:** The uplink ports on the QFX5120-48Y switches support auto-channelization.

  If you have disabled auto-channelization, then to channelize the ports, manually configure the port speed using the **set chassis fpc slot-number port port-number channel-speed speed** command, where the speed can be set to 10G or 25G. If a 100-Gigabit Ethernet transceiver is connected, you can only set

the speed to 25G. For the SFP+ ports, you can set the speed to 25G or 1G. There is no commit check for this, however.

> **NOTE:** You cannot configure channelized interfaces to operate as Virtual Chassis ports.

See [Channelizing Interfaces on Switches].

- **Resilient hashing support for link aggregation groups and equal-cost multipath routes (QFX5120 switches)**—Starting with Junos OS Release 18.3R1, resilient hashing is supported by link aggregation groups (LAGs) and equal-cost multipath (ECMP) sets on QFX5120 switches. A LAG combines Ethernet interfaces (members) to form a logical point-to-point link that increases bandwidth, provides reliability, and allows load balancing. Resilient hashing enhances LAGs by minimizing destination remapping when a new member is added to or deleted from the LAG. Resilient hashing works in conjunction with the default static hashing algorithm. It distributes traffic across all members of a LAG by tracking the flow's LAG member utilization. When a flow is affected by a LAG member change, the Packet Forwarding Engine rebalances the flow by reprogramming the flow set table. Destination paths are remapped when a new member is added to or existing members are deleted from a LAG. Resilient hashing applies only to unicast traffic and supports a maximum of 1024 LAGs, with each group having a maximum of 256 members. An ECMP group for a route contains multiple next-hop equal cost addresses for the same destination in the routing table. (Routes of equal cost have the same preference and metric values.) Junos OS uses a hash algorithm to choose one of the next-hop addresses in the ECMP group to install in the forwarding table. Flows to the destination are rebalanced using resilient hashing. Resilient hashing enhances ECMPs by minimizing destination remapping when a new member is added to or deleted from the ECMP group.

  [See Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups.]

Junos Telemetry Interface

- **Routing Engine and Packet Forwarding Engine sensors for the Junos Telemetry Interface (EX4650 and QFX5120-48Y switches)**—Starting with Junos OS Release 18.3R1, Routing Engine and Packet Forwarding Engine statistics are supported through the Junos Telemetry Interface on EX4650 and QFX5120-48Y switches with the same level of support found on QFX5100 switches using Junos OS Release 18.1R1.

  The following Routing Engine statistics are supported through JTI:

  - LACP state export
  - Chassis environmentals export
  - Network discovery chassis and components
  - LLDP export and LLDP model
  - BGP peer information (RPD)
  - RSVP interface export

- RPD task memory utilization export
- LSP event export
- Network Discovery ARP table state
- Network Discovery NDP table state

The following Packet Forwarding Engine statistics are supported through JTI:

- Congestion and latency monitoring
- Logical interface
- Filter
- Physical interface
- LSP
- NPU/LC memory
- Network Discovery NDP table state

Only gRPC streaming is supported.

To provision the sensor to export data through remote procedure call (gRPC), use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires you to download the OpenConfig for Junos OS module.

[See Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **Expanded physical interface queue and traffic statistics sensors for Junos Telemetry Interface (JTI) (PTX, MX, EX, QFX, ACX)**—Starting with Junos OS Release 18.3R1, additional resource paths are added to stream physical (IFD) statistics.

  Prior to Junos OS Release 18.3R1, both traffic and queue statistics for physical interfaces (IFD) are sent out together using the resource path **/interfaces** for gRPC streaming (which is internally used to create **/junos/system/linecard/interface/**) or **/junos/system/linecard/interface/** for UDP (native) sensors.

  Now, traffic and queue statistics can be delivered separately. Doing so can reduce the reap time for non-queue data for platforms supporting Virtual Output Queues (VOQ).

  The following UDP resource paths can be configured:

  - **/junos/system/linecard/interface/** is the existing resource path (no change). Traffic and queue statistics are sent together.
  - **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
  - **/junos/system/linecard/interface/queue/** exports queue statistics.

  The gRPC resource path **/interfaces** now has the following behavior:

  - In releases prior to Junos OS 18.3R1, it delivers all IFD traffic and queue statistics. In Junos OS 18.3R1 and higher, it delivers statistics in two sensors:

- **/junos/system/linecard/interface/traffic/** exports all fields except queue statistics.
- **/junos/system/linecard/interface/queue/** exports queue statistics.

To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. For streaming through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

[See sensor (Junos Telemetry Interface), Configuring a Junos Telemetry Interface Sensor (CLI Procedure), and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

For exporting statistics using UDP native sensors, configure parameters at the **[edit services analytics]** hierarchy level.

Layer 2 Features

- **Layer 2 unicast features(QFX5120 switches)**—Starting with Junos OS Release 18.3R1, the following Layer 2 unicast features are supported:
  - 802.1Q VLAN trunking
  - PVLAN
  - IRB
  - Layer 3 Vlan-tagged logical interfaces
  - 4096 VLANs
  - MAC address filtering
  - MAC address aging configuration
  - Static MAC address assignment for interface
  - Per-VLAN MAC learning (limit)
  - MAC learning disable
  - Persistent MAC (sticky MAC)
  - Q-in-Q Tag manipulation
  - MAC address limit per port
  - MAC limiting
  - MAC limiting per port, per VLAN
  - MAC move limiting
  - PVLAN on Q-in-Q
  - 802.1D

- 802.1w (RSTP)
- 802.1s (MST)
- BPDU protection
- Loop protection
- Root protection
- VSTP
- RSTP and VSTP running concurrently
- Link aggregation (static and dynamic) with LACP (fast and slow LACP)
- LLDP
- Multiple VLAN Registration Protocol (802.1ak)

See Ethernet Switching User Guide.

- **Layer 2 unicast features (QFX5120 switches)**—Starting with Junos OS Release 18.3R1, you can use the unified forwarding table (UFT) feature to allocate forwarding table resources to optimize the memory available for different address types based on the needs of your network. You can choose to allocate a higher percentage of memory for one type of address or another.

  [See Understanding the Unified Forwarding Table.]

MPLS

- **Support for MPLS over UDP tunnels (QFX10000 switches)**— Starting with Junos OS Release 18.3R1, MPLS-over-UDP tunnels are supported on QFX10000 switches. For every dynamic tunnel configured on the switch a tunnel composite next hop, an indirect next hop, and a forwarding next hop is created to resolve the tunnel destination route. You can also use policy control to resolve the dynamic tunnel over select prefixes by including the **forwarding-rib** configuration statement at the **[edit routing-options dynamic-tunnels]** hierarchy level.

  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.

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

- **MPLS support (QFX5120)**—Starting with Junos OS Release 18.3R1, the following MPLS features are supported:
  - LDP (tunneling over RSVP, targeted LDP, LDP over RSVP)
  - RSVP-TE
  - TE++ container LSPs
  - Automatic bandwidth allocation on LSPs
  - IPv6 tunneling over an MPLS IPv4 network (6PE and 6VPE)

- Ethernet-over-MPLS (L2 circuit)
- Layer 3 VPN
- Carrier-of-carrier VPNs
- ECMP routing
- Segment routing
- EVPN-VXLAN
- MPLS over IRB interfaces
- VRF support in IRB Interfaces

[See MPLS Feature Support on QFX Series and EX4600 Switches.]

Multicast

- **Layer 3 multicast features (QFX5120)**—Starting with Junos OS Release 18.3R1, the following Layer 3 multicast features are supported:
  - IGMP version 1 (IGMPv1), version 2 (IGMPv2), and version 3 (IGMPv3)
  - IGMP filtering
  - PIM sparse mode (PIM-SM)
  - PIM dense mode (PIM-DM)
  - PIM source-specific multicast (PIM-SSM)
  - Multicast Source Discovery Protocol (MSDP)

  IGMP and PIM are also supported on virtual routers.

  [See Multicast Overview.]
- **Layer 2 multicast features (QFX5120)**—Starting with Junos OS Release 18.3R1, the following Layer 2 multicast features are supported:
  - IGMP snooping for IGMPv1, IGMPv2, and IGMPv3
  - IGMP proxy
  - IGMP querier

  IGMP snooping is also supported on virtual routers.

  [See Multicast Overview.]

Network Management and Monitoring

- **Customized MIBs for sending custom traps based on syslog events (QFX Series)**—Starting in Junos OS Release 18.3R1, there is a process whereby customers can define their own MIBs for trap notifications. The customized MIB maps a particular error message with a custom OID rather than a generic one.

Juniper Networks provides two new MIB roots reserved for customer MIBs, one for the custom MIB modules and the other for the trap notifications. For this process, you must convert the MIB to YANG format, and a tool is available for that.

[See Customized SNMP MIBs for Syslog Traps.]

- **Services support: sFlow, port mirroring, and storm control (QFX5120 switches)**—Starting in Junos OS Release 18.3R1, the following services are provided on QFX5120 switches:
  - sFlow networking monitoring technology—Collects samples of network packets and sends them in a UDP datagram to a monitoring station called a *collector*. You can configure sFlow technology on a device to monitor traffic continuously at wire speed on all interfaces simultaneously.
  - Local and remote port mirroring and remote port mirroring to an IP address—Copies packets entering or exiting a port or entering a VLAN and sends the copies to a local interface (local port mirroring), to a VLAN (remote port mirroring), or to the IP address of a device running an analyzer application on a remote network (remote port mirroring to an IP address [GRE encapsulation]). (When you use remote port mirroring to an IP address, the mirrored packets are GRE-encapsulated.)
  - Storm control—Causes a device to monitor traffic levels and take a specified action when a specified traffic level—called the *storm control level*—is exceeded, thus preventing packets from proliferating and degrading service. You can configure devices to drop broadcast and unknown unicast packets, shut down interfaces, or temporarily disable interfaces when the storm control level is exceeded.

  [See Overview of sFlow Technology, Understanding Port Mirroring, and Understanding Storm Control.]

- **New fallback option for sFlow adaptive sampling (QFX Series)** —Starting with Junos OS Release 18.3R1, you can use the new CLI option *adaptive-sampling fallback* in sFlow monitoring configurations to back up the adaptive sampling rate on switch interfaces.

  Currently, adaptive sampling uses a binary *backoff* algorithm to reduce the sampling loads on the selected interfaces. However, if the sampling rate suddenly increases because of a spike in traffic, it does not revert to the previously configured value even after traffic stabilizes. *Adaptive sampling fallback* uses a binary *backup* algorithm to back up and decrease the sampling rate without affecting normal traffic. To enable this feature, include the **adaptive-sample-rate fallback** statement at the [**edit protocols sFlow**] hierarchy level. Adaptive sampling fallback is disabled by default. [See Understanding How to Use sFlow Technology for Network Monitoring.]

Restoration Procedures and Failure Handling

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (QFX Series)**—Starting in Junos OS Release 16.1, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, there is an automatic device recovery mode that goes into action should the system go into amnesiac mode.The new process is for the system to automatically retry to boot with the saved rescue configuration. In this circumstance, the system displays a banner "Device is in recovery mode" in the CLI (in both the operational and configuration modes). Previously, there was no automatic process to recover from amnesiac mode. A user with load and commit permission had to log in using the console and fix the issue in the configuration before the system would reboot.

  [See Saving a Rescue Configuration File.]

Routing Protocols

- **Junos OS, OpenConfig, and Network Agent packages are delivered in a single TAR file (QFX Series)**—Starting in Junos OS Release 18.3R1, the Junos OS image includes the OpenConfig package and Network Agent; therefore, you do not need to install OpenConfig or Network Agent separately on your device.

  [See Installing the OpenConfig Package and Installing the Agent Network Package.]

Security

- **Support for firewall filters (QFX5120)**—Starting with Junos OS Release 18.3R1, you can configure firewall rules to filter incoming network traffic based on a series of user-defined rules. You can specify whether to accept, permit, deny, or forward a packet before it enters an interface. If a packet is accepted, you can also configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). You configure firewall filters at the **[edit firewall]** hierarchy level.

  [See Firewall Filters Overview.]

- **Support for distributed denial-of-service protection (QFX5120)**—Starting with Junos OS Release 18.3R1, you can configure denial-of-service (DoS) protection on the switches to continue to function while under attack. A denial-of-service (DoS) attack is any attempt to deny valid users access to network or server resources by using up all the resources of the network element or server. DDoS protection identifies and suppresses malicious control packets while enabling legitimate control traffic to be processed. A single point of DDoS protection management enables you to customize profiles for your network control traffic. To protect against DDoS attacks, you can configure policers for host-bound exception traffic. The policers specify rate limits for all control traffic for a given protocol. You can also monitor policers, obtaining information such as the number of violations encountered and the number of packets received or dropped.

  [See Understanding Distributed Denial-of-Service Protection on QFX Series Switches.]

System Management

- **Secure boot (QFX5120 switches)**—Starting with Junos OS Release 18.3R1, a significant system security enhancement is introduced: secure bBoot. The secure boot implementation is based on the UEFI 2.4 standard. The BIOS has been hardened and serves as a core root of trust. The BIOS updates, the bootloader, and the kernel are cryptographically protected. No action is required to implement secure boot.

User Interface and Configuration

- **Support for configuring the ephemeral database using the NETCONF and Junos XML protocols (QFX Series)**—Starting in Junos OS Release 18.3R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database. The ephemeral database provides a fast programmatic interface that enables multiple clients to simultaneously load and commit configuration changes on a device running Junos OS and with significantly greater throughput than when committing data to the candidate configuration database. The device's active configuration is a merged view of the committed configuration database and the configuration data in all instances of the ephemeral configuration database. Ephemeral configuration data is volatile and is deleted upon rebooting the device.

  [See Understanding the Ephemeral Configuration Database.]

SEE ALSO

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for the QFX Series.

Interfaces and Chassis

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (QFX Series)**—In Junos OS Release 18.3R2, the **show lacp interfaces | display xml** command displays a new XML tag element **<lacp-hold-up-state>**. The **<lacp-hold-up-state>** displays the time interval an interface holds before it changes from state, down to up. In earlier Junos OS releases, the LACP hold up the information for all interfaces were in a single **<lacp-hold-up-information>** XML tag. Now, for each interface it is displayed in a separate **<lacp-hold-up-information>** XML tag.
- **Commit Error thrown when GRE interface and Tunnel source interface configured in different routing instances (QFX Series)**—In Junos OS Releases 18.3R2, QFX Series switches does not support configuring GRE interface and the underlying tunnel source interface in two different routing instances. If you try this configuration, it will result in a commit error with the following error message:

  error: GRE interface (gr-0/0/0.0) and its underlying tunnel source interface are in different routing-instances

  error: configuration check-out failed

  [See Understanding Generic Routing Encapsulation .]

Junos OS XML API and Scripting

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (QFX Series)**—Starting in Junos OS Release 18.3R1, Junos OS does not support configuring an MD5 or SHA-1 checksum hash to verify the integrity of local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) scripts or support using an MD5 or SHA-1 checksum hash with the **op url *url* key** option to verify the integrity of remote op scripts.

Network Management and Monitoring

- **Junos OS does not support management of YANG packages in configuration mode (QFX Series)**—Starting in Junos OS Release 18.3R1, adding, deleting, or updating YANG packages using the **run** command in configuration mode is not supported.
- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (QFX Series)**—Starting in Junos OS Release 18.3R2, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

Routing Policy and Firewall Filters

- **Support for configuring the GTP-TEID field for GTP traffic (QFX5000 line of switches)**—Starting in Junos OS Release 17.3R3, 17.4R2, 18.1R2, 18.2R1, and 18.3R1, the **gtp-tunnel-endpoint-identifier** statement is supported to configure the hash calculation of IPv4 or IPv6 packets that are included in the GPRS tunneling protocol–tunnel endpoint identifier (GTP-TEID) field hash calculations. The **gtp-tunnel-endpoint-identifier** configuration statement is configured at the **[edit forwarding-options enhanced-hash-key family inet]** hierarchy level.

  In most of the cases, configuring **gtp-tunnel-endpoint-identifier** statement is sufficient for enabling GTP hashing. After enabling, if GTP hashing does not work, it is recommended to capture the packets using relevant tools and identify the offset value. As per standards, 0x32 is the default header offset value. But, due to some special patterns in the header, offset may vary to say 0x30, 0x28, and so on. In this cases, use **gtp-header-offset** statement to set a proper offset value. Once the header offset value is resolved, run **gtp-tunnel-endpoint-identifier** command for enabling GTP hashing successfully.

  [See gtp-tunnel-endpoint-identifier and gtp-header-offset.]

Security

- **Syslog or log action on firewall drops packets (QFX5000 switches)**—Starting in 18.3R2, if you configure a syslog or log action on an ingress firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.
- **Firewall warning message (QFX5000 switches)**—Starting in 18.3R2, a warning message is displayed whenever a firewall term includes log or syslog with the accept filter action.

Virtual Chassis

- **New configuration option to disable automatic Virtual Chassis port conversion (QFX5100 Virtual Chassis)**—Starting in Junos OS Release 18.3R1, you can use the **no-auto-conversion** statement at the **[edit virtual-chassis]** hierarchy level to disable automatic Virtual Chassis port (VCP) conversion in a QFX5100 Virtual Chassis. Automatic VCP conversion is enabled by default on these switches. When automatic VCP conversion is enabled, if you connect a new member to a Virtual Chassis or add a new link between two existing members in a Virtual Chassis, the ports on both sides of the link are automatically converted into VCPs when all of the following conditions are true:
  - LLDP is enabled on the interfaces for the members on both sides of the link. The two sides exchange LLDP packets to accomplish the port conversion.
  - The Virtual Chassis must be preprovisioned with the switches on both sides of the link already configured in the members list of the Virtual Chassis using the **set virtual-chassis member** command.
  - The ports on both ends of the link are supported as VCPs and are *not* already configured as VCPs.

Automatic VCP conversion is not needed when using default-configured VCPs on both sides of the link to interconnect two members. On both ends of the link, you can also manually configure network or uplink ports that are supported as VCPs, whether or not the automatic VCP conversion feature is enabled.

Deleting the **no-auto-conversion** statement from the configuration returns the Virtual Chassis to the default behavior, which reenables automatic VCP conversion.

SEE ALSO

- New and Changed Features | **231**
- Known Behavior | **249**
- Known Issues | **253**
- Resolved Issues | **259**
- Documentation Updates | **270**
- Migration, Upgrade, and Downgrade Instructions | **271**
- Product Compatibility | **285**

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 for the QFX Series.

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

Class of Service (CoS)

- On QFX5000 line switches (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210), if the CoS configurations are modified when egress traffic is shaped at very low rate (lesser than 50 Mbps), packets might get stuck in the MMU buffers permanently. It might cause ingress and egress traffic drops. As a workaround, when low-rate shapers (lesser than 50 Mbps) are applied on egress queues, deactivate shaping before any CoS modification or that ensure traffic is stopped before doing CoS modification. PR1367432

EVPN

- When a VLAN uses an IRB interface as the routing interface, the **vlan-id** parameter must be set to **none** to ensure proper traffic routing. This issue is platform-independent. PR1287557
- IRB MAC/IP information will be deleted from ethernet-switching arp/nd table when **no-arp-suppression** is configured. PR1394959

Layer 2 Features

- The **Targeted-broadcast forward-only** command does not broadcast the traffic. PR1359031
- For QFX5120 and EX4650 the switch might learn its own MAC address on the network interface if it is attached an IRB interface to a VLAN. As a result of the incorrect MAC learning, it might result in the incorrect forwarding in a MC-LAG scenario. PR1365942
- Host table overflow occurs and routes are not programmed when the host table utilization is over 68 percent in the lpm-profile UFT for QFX5120 switches. PR1376581

Platform and Infrastructure

- In QFX10000, if the memory utilization exceeds the scale limit, dcpfe crashes. PR1329243
- When the sFlow collector can be reached only through Routing Engine, large samples due to heavy traffic can cause the Routing Engine CPU to become busy. PR1332337
- Hardware watchdog does not work on QFX10008, QFX10002-60C, and PTX10002-60C. PR1343131
- This issue is specific to flexible VLAN-tagged interfaces and does not happen if the interface is in trunk mode with EVPN-VXLAN configuration. PR1345568
- 100G Ethernet interface goes down after configure and deleting the ethernet loopback config. PR1353734
- **DIRECTORY CORRUPTED I=149350 OWNER=0 MODE=40755** messages continuously printed in console during device boot up after power cycle of the device The error logs are coming from inside Junos VM. As soon as any disk write operation is initiated from inside the VM, it will be written on host

disk as well. However, if power cycle happens before disk write completes, this issue is bound to occur. PR1361094

- IFL statistics are not supported for L2 and AE interfaces, it is supported only for L3 interfaces (L3 interface should not be member of AE), please make sure you have only normal L3 interface. PR1361185
- Bi-directional optics channelization is not supported. PR1361891
- In QFX5000 switches when more than one interface is attached to an output VLAN for remote port mirroring, the traffic will be received by only one of the interfaces. PR1363358
- A few harmless error messages related to function rt_mesh_group_add_check() will be seen during reboot. PR1365049
- auto channelization not supported for 40GBASE-BXSR QSFP+40GE-LX4 QSFP-100G-PSM4 100GBASE-BXSR. PR1366103
- In QFX5120 there will be 50-60s delay in **show ethernet-switching table summary** output to show all MACs when 288K MAC learned. PR1367538
- Sub-second BFD interval timer is not supported for QFX5120 switches. PR1368671
- "Boot [P]revious installed Junos packages " is not working from OAM boot menu option. Since this is a VM based system, the recovery would be done from Linux recovery. PR1371014
- Subsecond BFD interval timer is not supported for QFX5120 and EX4650 switches. PR1368671
- For QFX5120 and EX4650 switches Vm based system the recovery would be done from LINUX recovery. PR1371014
- A bug in PTP-FPGA is causing all the streams to follow the announce rate of the first master stream created on the PTP FPGA, instead of the announce rate of the corresponding stream. As a result, all other downstream instances end up receiving the announce rate of the first stream, even though the negotiated rates are different for these streams. As a workaround is to configuring same announce rate for all the downstream instance. PR1383203
- Intermittently after Junos OS reboot two of channelized 25-Gigabit ports using 4x25-Gigabit breakout cable might not come up. PR1384898
- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. Device can be recovered using power-cycle of the device. PR1385970
- Re-ARP request is sent without VLAN ID, so RE-ARP fails. PR1390794

Routing Protocols

- Could scale IS-ISv4, 254 neighbor and 200k routes together. Beyond 200k routes with 254 neighbor, adjacency flaps and thus traffic drop is noticed. However, with 40 neighbor, 351k routes got scaled. PR1368106

- Since the flex counters are shared among IFPs and other tables, in an uni-dimensional testing, ipmc stats counter created will not be equivalent to number of ipmc entries created, and stat counter creation will fail with the error **No resources for operation** after 60000 entries. PR1371399
- The mcsnoopd error messages are seen in logs while adding or deleting IGMP PIM configuration. These are debug messages and are not harmful. PR1371662

User Interface and Configuration

- **Auto-complete caution for QFX10002-60c and PTX10002-60c personalities**—Starting in Junos OS Release 18.3R2, for QFX10002-60c and PTX10002-60c personalities, do not use auto-complete to display the list of arguments for the **request system software delete** command. You must look for the package name using the **show system software** command and then explicitly type the software package name in the **request system software delete** command.

  [See request system software delete].

Virtual Chassis

- A Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on a QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop(>2s) might occur. PR1347902

SEE ALSO

New and Changed Features | **231**

Changes in Behavior and Syntax | **245**

Known Issues | **253**

Resolved Issues | **259**

Documentation Updates | **270**

Migration, Upgrade, and Downgrade Instructions | **271**

Product Compatibility | **285**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 18.3R2.

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

EVPN

- Mac-move-shutdown stops working if a physical loop is introduced continuously in quick succession of 10 minutes. The issue is not seen every time but can occur only if the physical loop is introduced at least four times. If the loops span a long period, the issue is not seen. PR1284315
- Chained-composite-next-hop (CNH) is a mandatory requirement for EVPN pure type 5 with VXLAN encapsulation. Without this, the Packet Forwarding Engine does not program the tunnel next hop. You have to explicitly set it on QFX5110 using the **set routing-options forwarding-table chained-composite-next-hop ingress evpn** command. For QFX10000, it is applied as part of the default configuration. PR1303246
- In a EVPN collapsed Layer 2/Layer 3 multi-homed GWs topology, when traffic is sent from IP Fabric towards EVPN, some traffic loss is seen. If the number of hosts behind EVPN gateways are increased, the traffic loss becomes higher. This issue is seen with QFX10000 switches. PR1311773
- A core link flap might result in inconsistent global MAC count. PR1328956

- At times, when l2ald is restarted, a race condition occurs where VTEP notification comes in from the kernel before lo0. As a result, l2ald is unable to process the VTEP add request and gets stuck in an indefinite loop. PR1384022
- To filter and see the output of desired ESI or neighbor information of an EVPN instance, there are two new choices available: **show evpn instance <> esi-info esi <>** and **show evpn instance <> neighbor-info neighbor <>**. PR1402175

General Routing

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging has been enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. PR1075230
- Layer 3 multicast traffic does not converge to 100 percentage and continuous drops are observed after bringing down or bringing up the downstream interface or while an FPC comes online after an FPC restart. This happens with multicast replication for 1000 VLANs or IRBs. PR1161485
- Single-bit and multiple-bit ECC errors are not logged on QFX5110 switches. PR1251917
- PathErr messages are not being received on link failure after disabling the interface. PR1275392
- On the QFX10002 platform, SXE interfaces erroneously configured in configuration might cause MAC pause frames to be generated on these internal interfaces and cause a Packet Forwarding Engine lockup. As a workaround, delete SXE interfaces from the configuration and then reboot. PR1281123
- Traffic drop occurs on sending traffic over et- interfaces due to CRC errors. PR1313977
- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED stays unlit. PR1317750
- There might be a traffic loss on the ingress PE device if the EVPN MPLS is configured later on the remote PE device or from the working condition EVPN MPLS is disabled and enabled later. PR1319770
- On the QFX10002-60C, filter operation with log action is not supported for protocols other than Layer 2, IPv4, and IPv6. The following message is seen in firewall logs: **Protocol 0 not recognized**. PR1325437
- BFD session over aggregated Ethernet flaps when a member link carrying the BFD Tx flaps. PR1333307
- On QFX10002, QFX10008, and QFX10016, ND is incorrectly working on IRB/Layer 3 interface with discard filter. PR1338067
- On QFX5100 platforms with sFlow enabled, when deleting or deactivating the sFlow interface, all other interfaces might go down and fxpc core files will be generated. PR1356868
- When **MC-LAG** is configured with **force-up** enabled on MCLAG nodes, the LACP admin key should not match the key of the access or CE device. PR1362346
- When a VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter will not be installed. PR1362609

- QFX5120: After sFlow configuration deletion, for every 5 minutes the **sflow_net_socket_init, 423sflow socket connect failed (socket closed)** error message will be displayed on the VTY console. PR1363381
- On QFX5000 switches, if lcmd is restarted, a chassisd core file will be generated with traffic drop for few seconds. PR1363652
- QFX52100: Filter with the routing-instance applied to family inet logical interfaces causes traffic to be discarded on unrelated interfaces. PR1364020
- The time lapse between interface-down interrupt detection to FRR callback is approximately 148 ms on the QFX5120 switch, though the in-place update FRR programming completes in 1 ms. The minimum FRR time achieved with this limitation is approximately 150 ms and maximum is approximately 275 ms. PR1364244
- Force-host upgrade is required for QFX5110-48S-4C in Junos OS Release 18.3 if the PTP over IPv6 G.8275.2 feature is required. PR1364735
- On the QFX5200, an error might be encountered when upgrading from Junos OS Release 15.1X53-D230.3 (the image with enhanced automation support [flex]) to an Junos OS Release 18.1R1.9 image without the enhanced automation. PR1366080
- Dedicated minimum buffers are reserved for some queues according to the Junos OS working model. These buffers are always available to those queues irrespective of the traffic pattern throughout the system. When the **clearing stat** statement is used, these values are visible. This cosmetic or minor issue has no functional impact. PR1367978
- Immediately after the AIS script package is installed, if any CLI command is executed, then no output is generated. PR1368039
- Accton AS7816-64X systems are shipping with 14 characters but Junos limitation is 12 characters. Accton serial number contains 781664X as the first 7 characters and 78 should be added from the **show chassis hardware** command output when the serial number is required. PR1371126
- When there are a large scale of VLANs around 4000, and if you add and delete VLANs as part of the same commit or two different commits with less time interval between them, then the VLAN tokens in the kernel will be exhausted due to which some of the VLANs will not get tagged. The following error message is seen: **/kernel: dcf_ng_vlan_alloc_hw_token: Couldn't allocate hardware token 65535 err=1** PR1371445
- MAC learning does not happen after restarting the l2-learning daemon for interfaces on backup. Traffic still gets forwarded. PR1372220
- On QFX5100 and QFX5200 Series Virtual Chassis platforms with GRES configured, if the backup member has the **/var/run/consoleredirect.pid** file, then rebooting the master member or while performing a Router Engine switchover, the backup cannot become the master member. PR1372521
- In Junos OS Release 18.1R3, when one 50-Gigabit Ethernet port is taken down using the **ifconfig** command, the other one also goes down. PR1376389
- LOC and Diag System LEDs on the front panel are not defined yet. PR1380459

- Last reboot reason is not correct if the device is rebooted because of power cycle. Last reboot reason will be displayed as Vjunos reboot even if the device got rebooted due to power cycling. PR1383693
- 1. Ingress VLAN-based mirroring is supported only using analyzer statement and does not work with firewall-based configuration. 2. Ingress VLAN mirroring is not supported with other firewall filters using VLAN on which VXLAN is enabled as match condition. 3. Ingress VLAN mirroring has to be configured again if the VLANs are deleted or if the EVPN-VXLAN configuration is deleted. PR1384732
- On QFX10008 and QFX10016 switches, traffic loss might be observed because of switch modular failure on the Control Board (CB). This failure further causes all SIBs to be marked as faulty and causes FPCs to restart until Routing Engine switchover occurs. PR1384870
- Junos OS might hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. You can recover the device by power cycling it. PR1385970
- In rare cases, rpd could end up with stuck KRT queue entries (visible in the **show krt queue** command output) as a result of interface flaps when using VRF configurations along with a static default route to em0.0 interface. PR1386475
- When **show** command is takes a long time to display results, the STP might changes states as BPDUs are no longer processed and cause outages. PR1390330
- On QFX10000 switches, the major alarm **FPC Management Ethernet Link Down** might be displayed for management Ethernet (em0 or em1) interface that is administratively down. The alarm message has no service impact and can be ignored. PR1391949
- DC Packet Forwarding Engine does not come up in some instances of abrupt power-off on power-on of QFX5120 switches. Power cycle the device or perform a host reboot to recover the device. PR1393554
- Changing of VNI underlay is not supported. PR1397999
- On QFX5100 switches, traffic initiated from a server connected to an interface will be dropped at the interface on the switch if the interface is configured with family ethernet-switching with VXLAN and the configuration is changed to family inet. PR1399733
- QFX5120: OVSDB-managed VXLAN sees traffic loss. PR1401943
- Transition from collapsed to non-collapsed L2/L3 GW and vice versa needs switch reload due to stale source VTEP IP. PR1405956
- On QFX10002, QFX10008, and QFX10016 switches, an auto-correctable non-fatal hardware error on PE chip (which is ASIC on PTX1000, PTX10002, QFX10002, the third-generation FPC on PTX3000/PTX5000, and the line card on PTX10008, PTX10016, QFX10008, and QFX10016) is reported as a 'FATAL' error and hence the related Packet Forwarding Engine will be disabled. The code changes have been made to change the error category from 'FATAL' to 'INFO' to avoid the Packet Forwarding Engine getting disabled unexpectedly. PR1408012
- On QFX5120 platforms with QSFP-100G-PSM4 transceiver, due to a timing fault on Field Programmable Gate Array (FPGA) hardware, the link might go down due to TX laser being disabled. PR1410687
- On QFX 5110 and QFX 5120switches, uRPF check in strict mode might not work properly. PR1417546

- During repeated power cycle tests, occasionally it is observed that 100-Gigabit PSM4 optics go to a state where link do not come up. This issue occurs frequently in a negative temperature environment (below -5 degrees Celsius). Physically reseating the transceiver or power cycling the device will help recover from issue state. PR1419826
- For transit static LSPs, QFX5120-32C devices might end up in swapping with an invalid label instead of POP/PHP action and might result in packet drop in the adjacent LER node. Since TD3 chipset has additional capabilities for MPLS, this issue is applicable only to QFX5120-32C platforms and not applicable to other platforms. As a workaround, removing/re-applying the static transit LSP configuration will solve this issue. PR1420370
- On QFX10002/QFX10008/QFX100016 platform, if BFD session if configured on fast mode, when the BFD session is across a dual-tagged Layer 3 interfaces (for example QinQ), BFD might stuck in slow mode. PR1422789
- MCLAG MAC synchronization not happening for local MACs when a new primary Lo0 IP is added and removed. This issue ocurrs when there are two primary addresses configured on the Lo0 and one is removed later. In this case, all IFBDs are removed and added back due to static vtep configuration. RG-ID of BD is reset and MAC synchronization does not happen. This happens for few VLANs only. PR1424013
- On QFX10000 platforms (QFX10002, QFX10008, QFX10016), if BFD is configured, heap memory leak might be seen. PR1427090

Infrastructure

- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. PR1315605

Layer 2 Features

- The **Targeted-broadcast forward-only** command does not broadcast traffic. PR1359031
- Trraffic is dropped when the core-side interface is configured as an IRB interface. PR1394952
- On QFX/ACX5000, for interfaces where LLDP is already disabled (commit) and there is any change on any interface in the next commit, l2cpd sends the message to disable LLDP on the all the interfaces to kernel. The kernel tries to remove the implicit filters, which return ENOENT, since entries were already disabled during the first commit. PR1400606
- After upgrading QFX5100 to Junos OS Release 14.1X53-D48, storm control does not take effect despite the profile taking effect. PR1401086

MPLS

- There could be some lingering RSVP state that would keep some labeled routes programmed in the Packet Forwarding Engine for a longer time than they should be. This RSVP state will eventually expire and then delete the RSVP MPLS routes from the FIB. However, traffic loss is not anticipated due to this lingering state or the corresponding labeled routes in the FIB. In the worst case, in a network, where there is persistent link flapping going on, this lingering state could interfere with the LSP scale being achieved. PR1331976
- Statistics of transit traffic does not increment LSP statistics signaled by RSVP-TE. PR1362936

Platform and Infrastructure

- When chassis control restart is done with the COS rewrite rule configured on aggregated Ethernet interface, the **Platform failed to bind rewrite** messages might be seen in syslog. The issue is specific to aggregated Ethernet interfaces. It is a timing issue that can occur when logical interface deletion is delayed due to high scale. When the logical interfaces come up again after restart, they have different indices. PR1315437

Routing Protocols

- We strongly recommend using BGP as the protocol for configuring the local address for each multihop iBGP/eBGP peer configuration. We recommend that local address be a route-able lo0 address. Using loopback address reduces dependency with interfaces. Note: multihop is by default enabled for iBGP peers. PR1323557
- Higher convergence time is observed for LFA with BFD in Junos OS Release 18.1. PR1337412
- In an MC-LAG setup, when status-control standby is rebooting and status-control active is down, and if ICCP session-establishment timer is configured less than or equal to the init-delay-timer on status-control standby, then mcae status of status-control standby might not become as active until the peer node is up. To avoid this issue,you should configure the ICCP session-establishment time to be greater than init-delay-timer with preferably 100 seconds or more. PR1348648
- On a scaled setup, when the host table is full and the host entries are installed in the LPM table, OSPF sessions might take more time to come up. PR1358289
- On a QFX5120/EX4650 with UFT configuration **num-65-127-prefix-4**, when the configuration is scaled with more than 64 prefix IPv6 routes, the **show Packet Forwarding Engine route inet6 hw lpm** command output will show only a single IPv6 entry but not the scaled entries. PR1369320
- On QFX Series switches except for QFX10000, if host destined packets (i.e., the destination address belongs to the device) come from the interface with ingress filter of log/syslog action (e.g., 'filter <> term <> then log/syslog'), such packets might not be dropped and reach the Routing Engine unexpectedly. PR1379718

- The **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(),128:l3 nh 6594 unintsall failed in h/w with Mini-PDT base configurations** error is seen on QFX5100 Virtual Chassis. There is no functionality impact due to this error message. PR1407175
- Inter-VLAN traffic duplication will be seen for some VLANs, after the configurations are loaded. PR1407920
- On QFX5110/QFX5200 platforms, the dcpfe might crash if any interface flaps. PR1415297

SEE ALSO

New and Changed Features | **231**

Changes in Behavior and Syntax | **245**

Known Behavior | **249**

Resolved Issues | **259**

Documentation Updates | **270**

Migration, Upgrade, and Downgrade Instructions | **271**

Product Compatibility | **285**

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.3R2

EVPN

- A few minutes of traffic loss might be observed during recovery from link failure. PR1396597
- VNI is not updated on default route 0.0.0.0/0 advertised by EVPN type 5 prefix when local configuration is changed. PR1396915
- In the non-collapsed (centralized) topology, when one of the two spines deactivates the underlay protocol (ospf), the leaf still points the virtual gateway MAC's next hop to the spine that is down PR1403524
- The rpd might crash after NSR switchover in an EVPN scenario. PR1408749

General Routing

- The 1-Gigabit copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. PR1286709
- Status LED on the chassis does not show up on QFX10002-60C. PR1332991
- FEC is incorrectly displayed on QFX10002 and QFX5110. PR1360948
- On QFX5110 with Junos OS Release 17.3R1, the following log messages are seen: **kernel: tcp_timer_keep: Dropping socket connection**. PR1363186
- Extended traffic loss might be observed when unified ISSU is performed with aggregated Ethernet interface configured with LACP protocol. PR1365316
- SFP-T might not work on QFX5100/QFX5110 devices. PR1366218
- For releases later than Junos OS Release 18.1R1, USB image installation on QFX5210-64C requires an AMI Bios upgrade. PR1371199
- The Packet Forwarding Engine is in a bad state after performing optics insertion or removal on a port. PR1372041
- The IPv6 routed packet might be transmitted through an interface whose VRRP state is in non-master. PR1372163
- MAC refresh packet might not be sent out from the new primary link after RTG failover. PR1372999
- On the QFX5110, the Ethernet switching flood group shows incorrect information. PR1374436
- RIPv2 update packets might not sent with IGMP snooping enabled. PR1375332
- A Packet Forwarding Engine wedge might be observed if there are interfaces going to down state. PR1376366
- Same address family [Subnet logical interface or IRB logical interface but not both] needs to be configured for establishing VTEPs. PR1376996
- The autonegotiation interface might go down if the opposite device supports only 10/100M autonegotiation. PR1377298

- The **expr_nh_flabel_check_overwrite: Caller nh_id params** debug log message is classified as error log when it should be LOG_INFO. PR1377447
- Deleting an IRB interface might affect other IRB interfaces if the same custom MAC address is configured. PR1379002
- The overlay ECMP might not work as expected on QFX5110 in an EVPN-VXLAN environment. PR1380084
- The Packet Forwarding Engine on QFX5000 might have DISCARD next hop for overlay-bgp-lo0-ip in the VXLAN scenario. PR1380795
- Traffic might be discarded without notification caused by FPC offline in a MC-LAG scenario. PR1381446
- The 40G-SR4 transceiver might not be recognized after a Junos OS upgrade on QF5100e. PR1381545
- SSD lifetime might be shortened in OVSDB environment. PR1381888
- LACP stuck in detached or attached state when an interface configured with native VLAN ID and VXLAN VLAN. PR1382209
- EVPN-VXLAN ARP/NDP proxy is not working. PR1382483
- The Packet Forwarding Engine might crash if the GRE destination IP is resolved over another GRE tunnel. PR1382727
- The **RPD_KRT_Q_RETRIES: list nexthop ADD: No such file or directory** log might be continuously shown after the rpd restarts. PR1383426
- DMA failure errors might be seen when the cache flushes or the cache is full. PR1383608
- The Virtual Chassis could not come up after upgrading to QFX5E platforms (TVP-based platforms for QFX5100 or QFX5200 switches). PR1383876
- The Layer 3 interface might stop pinging the directly connected link address after deleting Layer 2 on the physical interface. PR1384144
- On QFX5110platforms,SFPP-10G-DT-ZRC2 and SFPP-10G-CT50-ZR transceivers might not be tunable and remain 1550.10nmby default in the hardware. PR1384524
- Vm core file might be seen on the Junos OS Release 18.1R3. PR1384750
- Occasionally two of the channelized 25-Gigabit ports using 4x25-Gigabit breakout cable will not come up after Junos OS reboot. PR1384898
- All 1-Gigabit SFP copper and 1-Gigabit fiber optic links remain up on QFX10008 after all SIBs/FPCs are offline. PR1385062
- The IPv6 packet might not be routed when the IPv6 packet is encapsulated over IPv4 GRE tunnel on QFX10000. PR1385723
- The spine EVPN routes might be stuck in a hidden state with next hop as unusable after FPC is offline in the spine. PR1386147
- DDOS statistics and logging is not working for internal queues such as Q42 and Q4. PR1387508

- Traffic drop might be seen on QFX10000 platform with EVPN-VXLAN configured. PR1387593
- QFX5100, QFX5110, QFX5200, and QFX5210 Virtual Chassis could not be formed normally. PR1387730
- CPSM daemon memory leak in seen on VM host. PR1387903
- Certain log messages might be observed on QFX platforms. PR1388479
- ARP received on SP-Style interface is not sent to all RVTEPs in case of QFX5100 Virtual Chassis. Normal BUM traffic works fine. PR1388811
- MAC learning might stop working on some LAG interfaces. PR1389411
- FPC might crash on QFX5100 and EX4600 platforms in a large-scale scenario. PR1389872
- The vmcore might be seen when routing changes are made on the peer spine in an EVPN-VXLAN scenario. PR1390573
- An incorrect error message might be seen when J-Flow sensors are configured with reporting rate less than 30 seconds. PR1390740
- sdk-vmmd might consistently write to the memory. PR1393044
- 10-Gigabit Ethernet copper link flapping might happen during TISSU operation of QFX5100-48T switches. PR1393628
- IPv6 next hop programming issue might be observed on QFX10000 switches. PR1393937
- L2ALD core files are seen when **l2-learning traceoptions** are enabled. PR1394380
- DRAM and buffer utilization fields are not correct for QFX10000 platforms. PR1394978
- PTP over Ethernet traffic could be dropped if IGMP and PTP TC are configured together. PR1395186
- Unable to install licenses automatically on QFX Series platforms. PR1395534
- The subscriber bindings might not be successful on QFX/EX platforms. PR1396470
- On QFX5110, fan LED turns amber randomly. PR1398349
- High jsd or na-grpcd CPU usage might be seen even though JET or JTI is not used. PR1398398
- The DHCPv6 relay packets are dropped when both the UDP source and destination ports are 547. PR1399067
- CPU hog might be observed on PTX/QFX10000 switches. PR1399369
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. PR1399683
- ARP refresh functionality might fail in an EVPN scenario. PR1399873
- SFP-LX10 does not work on QFX5110. PR1399878
- PEM I2C failure alarm might be showed incorrectly as failed. PR1400380
- MAC-limit with persistent MAC is not working after reboot PR1400507

- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- The authd might crash when issuing the **show network-access requests pending** command during the authd restart. PR1401249
- File permissions are changed for **/var/db/scripts** files after reboot PR1402852
- The STP does not work when aggregated interfaces number is "AE1000" or above in QFX5000 and "AE480" or above in other QFX and EX Series switches. PR1403338
- The VRRP VIP might not work when it is configured on the LAG interface. PR1404822
- ARP/ND will not be resolved in case of native VLAN ID configured for LAG access interface. PR1404895
- A commit warning is seen on QFX5100. PR1405138
- VXLAN transit traffic over tagged underlay Layer 3 interface gets dropped due to hardware limitation. PR1406282
- The ARP request might not be resolved successfully if the **arp-suppression** is enabled and **vlan-id-list** is configured on the spine node. PR1407059
- DHCP discover packets getting dropped over VXLAN tunnel on a pure Layer 2 VLAN when DHCP relay is enabled for other VLANs. PR1408161
- The FPC might crash and does not come up if interface number or next hop is set to maximum value under **vxlan-routing** on QFX platforms. PR1409949

Interfaces and Chassis

- Constant dcpfe process crash might be seen if using an unsupported GRE interface configuration. PR1369757

Junos Fusion Provider Edge

- BUM traffic might get dropped on peer Fusion Aggregation Device when the link between the Satellite device and the local aggregate device goes down. PR1384440

Junos Fusion Satellite Software

- Extended Port (EP) LAG might go down on the Satellite Devices (SDs) if the related Cascade Port (CP) that links to an Aggregation Device (AD) goes down. PR1397992

Layer 2 Features

- The dcpfe process might crash while changing MTU of physical ports for GRE. PR1384517
- The LACP might be in detached state when deleting **native-vlan-id** on the aggregated Ethernet interface with **flexible-vlan-tagging** configured. PR1385409
- The dcpfe core file might be observed when doing **restart routing** or BGP neighbors flaps when EVPN-TYPE 5 routes are present. PR1387360

- The IPv6 NS/NA packets coming from the remote VTEP are not getting forwarded to the local host. PR1387519
- The dcpfe process might crash after VXLAN overlay ping. PR1388103
- With IGMP snooping enabled on the leaf switches, multicast traffic is forwarded to VLAN/VNI which does not have an active receiver. PR1388888
- RTG MAC refresh packets will be sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. PR1389695
- Packets destined to 01:00:0c:cc:cc:cc are not forwarded on QFX10000. PR1389829
- DCPFE restarted at the _bcm_field_td_counter_last_hw_val_update routine after upgrading spine with latest image. PR1398251
- With native VLAN (160) configured and host on non-native plan(100 -tagged) ARP packets sent with wrong VNI. PR1400000
- The dc-pfe process crash might be observed during restart of the Packet Forwarding Engine or system with scaled EVPN-VXLAN configuration. PR1403305
- The IPv6 NS/NA packets received over VTEP from an ESI host are wrongly flooded back to the host. PR1405820
- With **arp-suppression** enabled, QFX5K/EX46 might not forward IPv6 Router Solicitations or advertisements packets. PR1414496

Layer 2 Ethernet Services

- After GRES switchover, LACP might be down on the peer device and can never been recovered automatically. PR1395943

Multiprotocol Label Switching (MPLS)

- LSP "statistics" and "auto-bandwidth" functionality might not take effect with single-hop LSPs. PR1390445

Network Management and Monitoring

- Log files might not get compressed during the upgrade. PR1414303

Platform and Infrastructure

- Traffic might be discarded with indirect next-hop and load balancing. PR1376057
- IPv6 ping might fail for spine node in an EVPN scenario. PR1380590
- IRB interface does not go down when master of Virtual Chassis is rebooted or halted. PR1381272

Routing Protocols

- The pfe process might crash and all interfaces might flap as a result. PR1369011
- The rpd process might crash after committing the configuration related to mapping-server-entry. PR1379558

- Host-destined packets with filter log action might reach the Routing Engine. PR1379718
- BUM packets might get looped if EVPN multihoming interface flaps. PR1387063
- If a QFX5100 device has a host route with ECMP (equal-cost multipath) next-hops and receives a better path with single next-hop then next-hop in hardware will not be changed. PR1387713
- A dcfpe core file is seen at **brcm_pkt_tx_flush, l2alm_mac_ip_timer_handle_expiry_event_loc** after a random event. PR1397205
- The rpd core file might be seen when Layer 2 VPN is used. PR1398685

Resolved Issues: 18.3R1

Class of Service (CoS)

- A DST IP 224/4 match condition is programmed in the hardware as 224/24 in loopback FF entry rep=0. PR1354377

EVPN

- On a QFX10000 line switch with EVPN-VXLAN, **jprds_dlu_alpha_add : 222 JPRDS_DLU_ALPHA KHT** addition failed.PR1258933
- Logical interfaces from the same physical port do not work if configured under the same VXLAN VLAN. PR1278761
- When a VLAN uses an IRB interface as the routing interface, the vlan-id parameter must be set to "none" to ensure proper traffic routing. This issue is platform independent. PR1287557
- In EVPN-VXLAN environments, BFD flaps cause VTEP flaps and then Packet Forwarding Engine process crashes. PR1339084
- On QFX10000 line platforms with a scaling EVPN-VXLAN configuration, rpd generates a core file. PR1339979
- In EVPN-VXLAN scenarios, traffic might get silently dropped or directed to interfaces that are down, but LACP is up. PR1343515
- Traffic loss might be seen on Layer 2 and Layer nodes in a multihomed EVPN scenario. PR1355165
- The QFX10000 might drop transited traffic coming from the MPLS network to VXLAN-EVPN. PR1360159
- Increased risk of routing crash with temporary impact on traffic occurs on QFX10000 or QFX5100 nodes with certain configuration changes or when clearing L2 or L3 learning information in a high-scale EVPN-VXLAN configuration environment. PR1365257
- OSPF sessions are not coming up between MX Series routers and QFX10000 line switches as ARP entries get deleted and added. PR1366860

- Proxy ARP might not work as expected in an EVPN environment. PR1368911
- On QFX10000 line switches, importing the default IPv6 route to VRF causes infinite entries to get created in the EVPN internal IP prefix database and become unstable. PR1369166

Infrastructure

- QFX5100: Enabling **mac-move-limit** stops ping on the flexible-vlan-tagging enabled interface. PR1357742

Interfaces and Chassis

- Packets might drop on ICL of the MC-LAG peer where MC-LAG is up. PR1345316
- If the C-VLAN range is 16, it might not pass traffic in a Q-in-Q scenario. PR1345994

Junos Fusion Provider Edge

- Ppmd crashes after changing the mode of EX4300 from standalone to SD. PR1375647

Junos Fusion Satellite Software

- AD failure (power off) in a DC fusion is causing complete or partial traffic loss for an extended period. PR1352167

Layer 2 Features

- Broadcast frames might be modified with the ethertype 0x8850. PR1343575
- On random initialization of QFX5100 the programming of storm control profile in missed within hardware on random interfaces. This is not visible over CLI and the configuration still shows intact. PR1354889
- LACP packets are getting dropped with **native-vlan-id** configured after reboot. PR1361054
- The dcpfe or fxpc process might crash on Packet Forwarding Engines with low memory when allocating huge memory. PR1362332
- QFX5000 Virtual Chassis acting as EVPN-VXLAN ARP proxy might cause ARP resolution to fail. PR1365699
- Hashing does not work for the IPv6 packet encapsulated in a VXLAN scenario. PR1368258
- When **native-vlan-id** is configured for an AE interface, the LACP session to multihomed server goes down. PR1369424
- A port might still work even if it is deleted from an AE interface. PR1372577
- DHCP Discover packets might be dropped if VXLAN is configured. PR1377521

MPLS

- RSVP sessions go down for ingress LSPs with **no-cspf** enabled. PR1339916
- LSP is not received by QFX5110. PR1351055

- NO-propogate-TTL acts on the MPLS swap operation. PR1366804
- LSP with **auto-bandwidth** enabled goes down during an HMC error condition. PR1374102

Network Management and Monitoring

- For QFX5110, the returned snmp values of ModuleTemperature-HighAlarmThreshold/LowAlarmThreshold/HighWarningThreshold is not as same as the one shown on CLI. PR1369030

Platform and Infrastructure

- On the QFX10016 EVPN-VXLAN scaled testbed, it takes up to 3 minutes for traffic to converge when configuration. PR1323042
- The GRE traffic is not de-encapsulated by the firewall filter. PR1325104
- CoS is incorrectly applied on the Packet Forwarding Engine, leading to egress traffic drop. PR1329141
- The etherStatsCRCAlignErrors counters might disappear in the SNMP tree. PR1329713
- On QFX10000 line platforms, DHCP relay/server is not working on a GRE interface. PR1331158
- EVPN-VXLAN: Delay Factor drops multicast traffic. PR1333069
- Ethernet frames with Ethernet type of 0x8922 might be modified at egress by QFX10000 line platforms. PR1334711
- The device uses the well-known ports as source port in VXLAN scenario. PR1335227
- AI-script does not get automatically reinstalled during a Junos OS upgrade on a next-generation Routing Engine. PR1337028
- The Delay Factor of an EVPN instance might flood all the ARP requests back to the Ethernet segment. PR1337275
- On QFX5100 platforms, LR4 QSFP can take up to 15 minutes to come up after VC reboot. PR1337340
- On the QFX10000 platforms, VRRP function does not work well when it is configured on logical interfaces. PR1338256
- The VXLAN traffic might not be transmitted correctly with the IRB interface as the underlay interface of the VTEP tunnel. PR1338586
- On QFX5000 line platforms, DDoS counters for OSPF might not increase. PR1339364
- Multicast traffic drop is seen if downstream IRB interfaces have snooping enabled. PR1340003
- On QFX5100, QFX5200, QFX5110, and EX4600 platforms, BPDU packets might get dropped and **bpdu-block-on-edge** might not work. PR1343330
- PAFXPC core files were seen when remote member ifd was referenced in the "show dcbcm ifd <ifd-name> on QFX5100 Platform configured in a Virtual Chassis. PR1343701

- On QFX10000 line platforms, in an EVPN-VXLAN with flexible-tag mode deployment, 100G interface statistics do not get updated for ingress traffic. PR1343746
- On any platforms supporting EVPN-VXLAN, any IRB-sourced packet might use the VRRP/virtual-gateway MAC address in the Ethernet header instead of the IRB MAC address. PR1344990
- On the QFX5100, the fan RPM fluctuates when the temperature sensor reaches its threshold. PR1345181
- The fxpc process might crash when removing all VXLAN configuration. PR1345231
- The backup Routing Engine might crash, causing vmcore to be generated on the master Routing Engine, master Routing Engine performance will not be affected. PR1346218
- Incorrect inner VLAN tag is sent from QFX10K platform with Q-in-Q configured on the Layer 3 logical interface. PR1346371
- On QFX10000 line platforms, syslog error messages might be seen in syslog after configuring multiple LAG interfaces under the sFlow protocol. PR1346493
- QFX5100-48T 10G interface might be autonegotiated at 100M speed instead of 10G. PR1347144
- On QFX5110-48S-4C platforms, part numbers and serial numbers are not displayed for any of the 10G optics/DAC connected. PR1347634
- Traffic in which the destination MAC matches the virtual gateway MAC might be silently dropped or discarded. PR1348659
- On the QFX10002-60C, vmhost might generate a core file right after a GR interface is configured. PR1348932
- The BGP session might flap after changing the extended-vni-list under EVPN hierarchy. PR1349600
- QFX5100 40G port has an interoperability issue with some other vendors. PR1349664
- The pfed process might consume high CPU resources if subscriber or interface statistics are used at a large scale. PR1351203
- Dcpfe process might crash on QFX10000 switches. PR1351503
- The GTP traffic might not be hashed correctly for the AE interface. PR1351518
- Telemetry traffic does not leave the local device when the telemetry server is reachable through a virtual router routing-instance. PR1352593
- QFX5100 ARP fails after the change interface MAC address is changed. PR1353241
- RPC output is not showing failure when running **request system software add** with software already staged. PR1353466
- On QFX5110 platforms, SFP-LX10 might stay in up or down state when connected. PR1353677
- Alarm errors might be seen during startup on QFX10000. PR1354582
- Untagged packets might not be forwarded through the trunk port. PR1355338

- A commit error is observed if the device is downgraded from Junos OS Release 18.2 or 18.3 release to Release 17.3R3. PR1355542
- On LX10 SFPs on QFX5110 platforms, autonegotiation is not in effect with new configurations. PR1355746
- EVPN-VXLAN: the VXLAN traffic might be lost in EVPN type 2 and type 5 scenario. PR1355773
- **Load averages** output under **show chassis routing-engine** shows **nan** periodically. PR1356676
- The device cannot match on user-vlan-id for tunnel-terminated packets. PR1358669
- The IGMP membership report packets might not be forwarded over an interface on QFX10000 line switches. PR1360137
- On QFX10000 line platforms, packets will be dropped when **virtual-gateway-address** is configured on an IRB interface associated with a non-VXLAN VLAN. PR1360646
- The GTP traffic might not be hashed correctly on the AE interface. PR1361379
- On QFX10000 line platforms, the **clear services accounting statistics inline-jflow fpc-slot** command does not work. PR1362396
- The QFX5100 Virtual Chassis is unable to connect to the management address through the vme interface. PR1362437
- Traffic might not be forwarded when the member link of the AE interface is added or deleted. PR1362653
- 1G interface might stop working when "auto-negotiation" is off by default. PR1362977
- OSPF might remain in init status after firmware upgrade loading the Junos OS Release 14.1X53-D47.4 image. PR1362996
- On QFX10008, QFX10016, PTX1000, PTX5000, PTX10008, and PTX10016 platforms, MPLS exp rewrite might not work for IPv6 and IPv4 traffic. PR1364391
- Root password recovery process does not work. PR1365740
- The tagged traffic is dropped in the untagged EVPN-VXLAN scenario. PR1366336
- On PTX10002, QFX10002-60C, and QFX10000-30C platforms, some interfaces do not come up during initialization after a reboot. PR1368203
- On QFX5100, QFX5110, and QFX5200 platforms, IS-IS adjacency goes down when MTU 9192 is configured. PR1368913
- The 'commit' or 'commit check' might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- The **ipv4-dscp** command is affecting the CoS treatment of PTPoIPv6 packets at the egress queue on PTP BC and PTP OSC with G.8275.2.enh profile. PR1371064
- On QFX10000 line platforms, before Junos OS Release 17.3R3 code, the maximum number of ESI logical interfaces was 4000 in the Packet Forwarding Engine. PR1371414

- Packet is dropped after the filter on the interface is deleted. PR1372957
- TPI-50840 BUM traffic received on QFX5110 is not flooded to all remote VTEPs. PR1373093
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. PR1373807
- LLDP might stop fully working between a QFX10000 and a non-Juniper device. PR1374321

Routing Protocols

- The **rpf-check-policy**statement does not work as expected. PR1336909
- On QFX5110 platforms, setting MTU on an L3 interface does not take effect. PR1345495
- On QFX10000 line platforms, NETCONF SSH TCP port 830 traffic is hitting host path or an unclassified queue. PR1345744
- On QFX5100 and EX4600 platforms, parity errors in the L3 IPv4 table in the Packet Forwarding Engine memory might cause traffic to be dropped or silently discarded. PR1364657

Virtual Chassis

- Traffic loop might be seen during network port to Virtual Chassis Port(VCP) port conversion. PR1346851

SEE ALSO

New and Changed Features | **231**

Changes in Behavior and Syntax | **245**

Known Behavior | **249**

Known Issues | **253**

Documentation Updates | **270**

Migration, Upgrade, and Downgrade Instructions | **271**

Product Compatibility | **285**

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 18.3R2.

SEE ALSO

New and Changed Features | **231**

Changes in Behavior and Syntax | **245**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

Upgrading Software on QFX Series Switches

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

If you are not familiar with the download and installation process, follow these steps:

1. In a browser, go to https://www.juniper.net/support/downloads/junos.html.

   The Junos Platforms Download Software page appears.

2. In the QFX Series section of the Junos Platforms Download Software page, select the QFX Series platform for which you want to download the software.

3. Select **18.3** in the Release pull-down list to the right of the Software tab on the Download Software page.

4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 18.3 release.

   An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

   A login screen appears.

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

7. Download the software to a local host.

8. Copy the software to the device or to your internal software distribution site.

9. Install the new jinstall package on the device.

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

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

   `user@host>` **request system software add *source*/jinstall-host-qfx-5-x86-64-18.3-R1.n-secure-signed.tgz reboot**

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

   - ***/pathname***—For a software package that is installed from a local directory on the switch.
   - For software packages that are downloaded and installed from a remote location:
     - **ftp://*hostname*/*pathname***

- **http://*hostname/pathname***
- **scp://*hostname/pathname*** (available only for Canada and U.S. version)

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

Rebooting occurs only if the upgrade is successful.

> **NOTE:** After you install a Junos OS Release 18.3 **jinstall** package, you can issue the **request system software rollback** command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-x.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

> **NOTE:** The QFX10002-60C switch supports only the 64-bit version of Junos OS.

> **NOTE:** If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add *<pathname><source>*** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-18.1R1.9.tgz
```

If the Install Package resides remotely from the switch, execute the **request vmhost software add *<pathname><source>*** command.

For example:

```
user@switch> request vmhost software add
ftp://ftpserver/directory/junos-vmhost-install-qfx-x86-64-18.1R1.9.tgz
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10002 Switches

NOTE: If you are upgrading from a version of software that does not have the FreeBSD 10 kernel (15.1X53-D30, for example), you will need to upgrade from Junos OS Release 15.1X53-D30 to Junos OS Release 15.1X53-D32. After you have installed Junos OS Release 15.1X53-D32, you can upgrade to Junos OS Release 15.1X53-D60 or Junos OS Release 18.3R2.

NOTE: On the switch, use the **force-host** option to force-install the latest version of the Host OS. However, by default, if the Host OS version is different from the one that is already installed on the switch, the latest version is installed without using the **force-host** option.

If the installation package resides locally on the switch, execute the **request system software add *<pathname><source>* reboot** command.

For example:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-f-x86-64-18.3R2.n-secure-signed.tgz reboot
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* reboot** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-f-x86-64-18.3R2.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

> **NOTE:** Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add *<pathname><source>*** command.

To install the software on re0:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* re0** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add
/var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add *<pathname><source>* re1** command.

For example:

```
user@switch> request system software add
ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the **show version** command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI **delete chassis redundancy** command when prompted. If GRES is enabled, it will be removed with the **redundancy** command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the **[edit routing-options]** hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

   ```
   user@switch> configure
   ```

3. Disable Routing Engine redundancy:

   ```
   user@switch# delete chassis redundancy
   ```

4. Disable nonstop-bridging:

   ```
   user@switch# delete protocols layer2-control nonstop-bridging
   ```

5. Save the configuration change on both Routing Engines:

   ```
   user@switch# commit synchronize
   ```

6. Exit the CLI configuration mode:

   ```
   user@switch# exit
   ```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the **request system software add** command:

   ```
   user@switch> request system software add validate
   /var/tmp/jinstall-host-qfx-10-f-x86-64-18.3R2.n-secure-signed.tgz
   ```

   For more information about the **request system software add** command, see the CLI Explorer.

9. Reboot the switch to start the new software using the **request system reboot** command:

   ```
   user@switch> request system reboot
   ```

   > **NOTE:** You must reboot the switch to load the new installation of Junos OS on the switch.
   >
   > To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the **request system software delete *<package-name>*** command. This is your last chance to stop the installation.

   All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

   While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the **show version** command to verify the version of the software installed.

    ```
    user@switch> show version
    ```

    Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

    For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

    ```
    user@switch> request chassis routing-engine master switch
    ```

    For more information about the **request chassis routing-engine master** command, see the CLI Explorer.

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

    ```
    user@switch> show chassis routing-engine
    Routing Engine status:
      Slot 0:
        Current state                  Backup
        Election priority              Master (default)
    Routing Engine status:
      Slot 1:
        Current state                  Master
        Election priority              Backup (default)
    ```

14. Install the new software package using the **request system software add** command:

    ```
    user@switch> request system software add validate
    /var/tmp/jinstall-host-qfx-10-f-x86-64-18.3R2.n-secure-signed.tgz
    ```

    For more information about the **request system software add** command, see the CLI Explorer.

15. Reboot the Routing Engine using the **request system reboot** command:

```
user@switch> request system reboot
```

> **NOTE:** You must reboot to load the new installation of Junos OS on the switch.
>
> To abort the installation, do not reboot your system. Instead, finish the installation and then issue the **request system software delete jinstall *<package-name>*** command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the **show version** command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the **request chassis routing-engine master** command, see the CLI Explorer.

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
```

```
Routing Engine status:
  Slot 0:
    Current state                  Master
    Election priority              Master (default)
outing Engine status:
  Slot 1:
    Current state                  Backup
    Election priority              Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- Preparing the Switch for Software Installation on page 282
- Upgrading the Software Using Unified ISSU on page 282

Preparing the Switch for Software Installation

Before you begin software installation using unified ISSU:

- Ensure that nonstop active routing (NSR), nonstop bridging (NSB), and graceful Routing Engine switchover (GRES) are enabled. NSB and GRES enable NSB-supported Layer 2 protocols to synchronize protocol information between the master and backup Routing Engines.

  To verify that nonstop active routing is enabled:

  NOTE: If nonstop active routing is enabled, then graceful Routing Engine switchover is enabled.

  ```
  user@switch> show task replication
          Stateful Replication: Enabled
          RE mode: Master
  ```

  If nonstop active routing is not enabled (**Stateful Replication** is **Disabled**), see *Configuring Nonstop Active Routing on Switches* for information about how to enable it.

- Enable nonstop bridging (NSB). See *Configuring Nonstop Bridging on Switches (CLI Procedure)* for information on how to enable it.
- (Optional) Back up the system software—Junos OS, the active configuration, and log files—on the switch to an external storage device with the **request system snapshot** command.

Upgrading the Software Using Unified ISSU

This procedure describes how to upgrade the software running on a standalone switch.

To upgrade the switch using unified ISSU:

1. Download the software package by following the procedure in the Downloading Software Files with a Browser section in *Installing Software Packages on QFX Series Devices*.
2. Copy the software package or packages to the switch. We recommend that you copy the file to the **/var/tmp** directory.
3. Log in to the console connection. Using a console connection allows you to monitor the progress of the upgrade.
4. Start the ISSU:
   - On the switch, enter:

     ```
     user@switch> request system software in-service-upgrade /var/tmp/package-name.tgz
     ```

     where ***package-name*.tgz** is, for example, **jinstall-host-qfx-10-f-x86-64-18.3R2.n-secure-signed.tgz**.

     NOTE: During the upgrade, you cannot access the Junos OS CLI.

     The switch displays status messages similar to the following messages as the upgrade executes:

     ```
     warning: Do NOT use /user during ISSU. Changes to /user during ISSU may get
     lost!
     ISSU: Validating Image
     ISSU: Preparing Backup RE
     Prepare for ISSU
     ISSU: Backup RE Prepare Done
     Extracting jinstall-host-qfx-5-f-x86-64-18.3R2.n-secure-signed.tgz ...
     Install jinstall-host-qfx-5-f-x86-64-18.3R2.n-secure-signed.tgz completed
     Spawning the backup RE
     Spawn backup RE, index 0 successful
     GRES in progress
     GRES done in 0 seconds
     Waiting for backup RE switchover ready
     GRES operational
     Copying home directories
     Copying home directories successful
     Initiating Chassis In-Service-Upgrade
     Chassis ISSU Started
     ISSU: Preparing Daemons
     ISSU: Daemons Ready for ISSU
     ```

```
ISSU: Starting Upgrade for FRUs
ISSU: FPC Warm Booting
ISSU: FPC Warm Booted
ISSU: Preparing for Switchover
ISSU: Ready for Switchover
Checking In-Service-Upgrade status
  Item           Status                  Reason
  FPC 0          Online (ISSU)
Send ISSU done to chassisd on backup RE
Chassis ISSU Completed
ISSU: IDLE
Initiate em0 device handoff
```

NOTE: A unified ISSU might stop, instead of abort, if the FPC is at the warm boot stage. Also, any links that go down and up will not be detected during a warm boot of the Packet Forwarding Engine (PFE).

NOTE: If the unified ISSU process stops, you can look at the log files to diagnose the problem. The log files are located at **/var/log/vjunos-log.tgz**.

5. Log in after the reboot of the switch completes. To verify that the software has been upgraded, enter the following command:

   ```
   user@switch> show version
   ```

6. Ensure that the resilient dual-root partitions feature operates correctly, by copying the new Junos OS image into the alternate root partitions of all of the switches:

   ```
   user@switch> request system snapshot slice alternate
   ```

   Resilient dual-root partitions allow the switch to boot transparently from the alternate root partition if the system fails to boot from the primary root partition.

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 17.1, 17.2 and 17.3 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

SEE ALSO

- New and Changed Features | 231
- Changes in Behavior and Syntax | 245
- Known Behavior | 249
- Known Issues | 253
- Resolved Issues | 259
- Documentation Updates | 270
- Product Compatibility | 285

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 285

Hardware Compatibility

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

To determine the features supported on QFX Series switches in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. Find Feature Explorer at https://apps.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 SRX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.3R2 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

You can also find these release notes on the Juniper Networks Junos OS Documentation webpage, located at https://www.juniper.net/documentation/product/en_US/junos-os.

> **NOTE:** The SRX5K-SPC3 Services Processing Card was introduced in Junos OS Service Release 18.2R1-S1 and is supported in all subsequent Junos OS Releases. The features and functionalities of the SRX5K-SPC3 card are supported in Junos OS Release 18.3R1. Going forward, future improvements for SRX5K-SPC3 will be included in upcoming Junos OS Maintenance Releases.

New and Changed Features

IN THIS SECTION

Release 18.3R2 New and Changed Features

There are no new features in Junos OS Release 18.3R2 for the SRX Series devices.

Release 18.3R1 New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 18.3R1 for the SRX Series devices.

Junos OS Release 18.3R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550M, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800. Most security features in this release were previously delivered in Junos OS for SRX Series “X” releases from 12.1X44 through 15.1X49-D150. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D150 are not available in 18.3 releases.

Application Security

- **Downloading the Junos OS application signature package from a proxy server (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, you can download the Junos OS application signature

package from a proxy server. You can download and install application signature package hosted on an external server when a Web proxy is already deployed on your device.

To download the signature package by using a proxy server, configure a profile with host and port details of the proxy server, and use the **set services application-identification download proxy-profile *profile-name*** command to connect to the external server through a specified proxy server.

The download retrieves the application signature package from the Juniper Networks security website https://signatures.juniper.net/cgi-bin/index.cgi.

[See Predefined Application Signatures for Application Identification.]

- **Elliptic Curve Digital Signature Algorithm (ECDSA) cipher support (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, ECDSA cipher suites are supported in SSL proxy for digital signing. ECDSA ciphers are based on Elliptic Curve Cryptography (ECC). ECDSA cipher suites are available with smaller keys, and provide faster and more secure cryptography across the Internet.

  SSL proxy supports only the ECC certificate with the Elliptic Prime Curve 256-bit (P-256).

  [See SSL Proxy Overview.]

- **URL category-based routing (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, advanced policy-based routing (APBR) feature is enhanced to include URL categories as match criteria in an APBR profile to enable URL category-based routing. URL categories are based on destination IP address, and the category identification is leveraged from Enhanced Web Filtering and local Web filtering results from UTM. APBR uses the details to match traffic and route the matching traffic to a specified next-hop device.

  URL category-based routing enables redirecting the traffic based on a specific website or a URL category to ensure that the Web traffic arrives at the appropriate destination.

  [See Advanced Policy-Based Routing.]

Authentication and Access

- **IPv6 support for configuring the JIMS server and filtering IP addresses (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, IPv6 addresses are supported to connect the Juniper Identity Management Service (JIMS) primary server and secondary server, in addition to existing IPv4 address support. Also, IPv6 addresses are supported to configure a filter based on IP addresses for the advanced query feature, in addition to existing IPv4 address support.

  [See Understanding the SRX Series Advanced Query Feature for Obtaining User Identity Information from JIMS.]

Authentication, Authorization and Accounting

- **Support for password change policy enhancement (SRX Series)**—Starting in Junos OS Release 18.3R1, the Junos password change policy for local user accounts is enhanced to comply with certain additional password policies. As part of the policy improvement, you can configure the following:

- **minimum-character-changes**—The number of characters by which the new password should be different from the existing password.
- **minimum-reuse**—The number of older passwords, which should not match the new password.

[See password.]

Flow-Based and Packet-Based Processing

- **Selective stateless packet forwarding (SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 18.3R1, selective stateless packet forwarding services are supported on SRX1500, SRX4100, SRX4200, and SRX4600 devices in addition to the existing support on SRX300, SRX320, SRX340, SRX345, and SRX550M devices. Using selective stateless packet forwarding services, the device is configured to provide packet-based processing for selected traffic based on the firewall filter input terms. The remaining traffic that is not filtered is processed using flow-based forwarding.

  Selective stateless packet forwarding is supported on the following protocols:

  - IPv4
  - MPLS
  - CCC-Ethernet switching cross-connects

  [See Understanding Selective Stateless Packet-Based Services and Example: Configuring Selective Stateless Packet-Based Services for Packet-Based to Flow-Based Forwarding.]

GPRS

- **GTP tunnel enhancements (SRX1500, SRX4100, SRX4200, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 18.3R1, GPRS tunneling protocol (GTP) is enhanced to update the GTP tunnel and session lifetime to avoid GTP tunnel timeout issues. Even if the GTP-U validation is disabled, the GTP-U traffic can refresh the GTP tunnel to avoid timeout. Only GTPv1 and GTPv2 tunnels, not GTPv0 tunnels, are refreshed by the GTP-U traffic. Before refreshing the GTP tunnel, you must enable the GTP-U distribution.

  > **NOTE:** On SRX5400, SRX5600, and SRX5800 devices, the number of GTP tunnels supported per SPU is increased from 200,000 tunnels to 600,000 tunnels, for a total of 2,400,000 tunnels per SPC2 card.

  [See Monitoring GTP Traffic.]

Intrusion Detection and Protection (IDP)

- **Downloading the IDP security package through an explicit proxy server (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, you can download the IDP security package through an explicit Web proxy server.

To download the IDP security package that hosts on an external server, you need to configure a proxy profile and use the proxy host and port details that are configured in the proxy profile.

This feature allows you to use a deployed Web proxy server on your device for access and authentication for HTTP and HTTPS outbound sessions.

[See Downloading the Junos OS IDP Signature Package through an Explicit Proxy Server.]

- **Support for multiple IDP policies (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, with unified policies configured on an SRX Series device, you can configure multiple IDP policies and set one of those policies as the default IDP policy. If multiple IDP policies are configured for a session and when policy conflict occurs, the device applies the default IDP policy for that session and thus resolves any policy conflicts.

  If you have configured two or more IDP policies in a unified security policy, then you must configure the default IDP policy.

  [See IDP Policies Overview.]

- **User visibility improvements for IDP attacks (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, you can view the attack objects that are available in an attack object group (predefined, dynamic, and custom attack groups) and the group to which a predefined attack object belongs.

  You can use the following new commands to view the details of attack objects in a group and the group to which a predefined attack belongs:

  - **show security idp attack attack-list attack-group *attack-group-name***
  - **show security idp attack group-list *attack-name***

  [See show security idp attack attack-list and show security idp attack group-list.]

Interfaces and Chassis

- **Management Ethernet interface (fxp0) is confined in a non-default virtual routing and forwarding table (SRX Series)**—Starting in Junos OS Release 18.3R1, you can confine the management interface in a dedicated management instance by setting a new CLI configuration statement, **management-instance**, at the **[edit system]** hierarchy level. By doing so, operators will ensure that management traffic no longer has to share a routing table (that is, the default.inet.0 table) with other control or protocol traffic in the system. Instead, there is a mgmt_junos routing instance introduced for management traffic.

  [See Management Interface in a Non-Default Instance and management-instance.]

Logical Systems and Tenant Systems

- **Application identification support enhancement for logical systems (SRX Series)**—Starting in Junos OS Release 18.3R1, the application identification (AppID) support for logical systems now includes two new options to display and clear logical system statistics and counters. The user logical system administrator can view the AppID signature package status and version. The custom signatures configured by the master logical system administrator can be configured in the user logical system security policies. You can view the information about AppID signature package status and version by using the commands **show services application-identification status** and **show services application-identification version**.

  [See Understanding Logical System Application Identification Services.]

- **ICAP redirect profile support for logical systems (SRX Series)**—Starting in Junos OS Release 18.3R1, SRX Series devices support the Internet Content Adaptation Protocol (ICAP) service redirect when the device is configured for logical systems.

  ICAP is a lightweight protocol used to extend transparent proxy servers, thereby freeing up resources. ICAP redirect profile is only allowed to attach on the policy that belongs to the same logical system.

  [See ICAP Redirects for Logical Systems.]

- **IDP support for logical systems (SRX Series)**—Starting in Junos OS Release 18.3R1, the intrusion detection and prevention (IDP) support is extended to logical systems.

  IDP support allows the following actions for logical systems:

  - Configure individual IDP policies.
  - Verify the IDP policy load and compilation status.
  - View the attacks detected and service statistics.

  A single IDP security package is installed at the master logical system that is shared by all other logical systems. Only the master logical system administrator can configure the **sensor-configuration** statement and this is used by other logical systems.

  [See IDP for Logical Systems.]

- **Logical systems support (SRX4600)**—Starting in Junos OS Release 18.3R1, SRX4600 device supports logical system in route mode only.

  [See Understanding Logical Systems for SRX Series Services Gateways.]

- **New context-oid option for trap-options configuration statement to distinguish the traps which come from a non-default routing instance and non-default logical system (SRX Series)**—In Junos OS Release 18.3R1, a new option, **context-oid**, for the **trap-options** statement allows you to handle prefixes such as <routing-instance name>@<trap-group> or <logical-system name>/<routing-instance name>@<trap-group> as an additional varbind.

  [See trap-options.]

- **Tenant systems support (SRX Series)**—Starting in Junos OS Release 18.3R1, tenant systems are supported. A tenant system provides logical partitioning of the SRX Series device into multiple domains similar to logical systems and provides high scalability. A tenant system supports routing, services and security features. A tenant system is created by the master administrator. The tenant system supports independent provisioning and administration. The master administrator uses the resource profiles to specify resource allocation for a tenant system. The tenant system administrator can configure and view the security features for the tenant systems.

  [See Tenant Systems Overview and Understanding Licenses for Logical Systems and Tenant Systems on SRX Series Devices.]

  The following features that are supported on the logical systems are now extended to tenant systems:

  - NAT (source NAT, destination NAT, and static NAT)

    [See NAT for Tenant Systems.]

  - Firewall authentication (pass-through authentication, Web authentication, user firewall authentication, and push authentication entries to Juniper Identity Management Service (JIMS))

    [See Firewall Authentication for Tenant Systems.]

  - ALG (data and VoIP)

    [See ALG for Tenant Systems.]

  - Security policies, zones, and logs

    [See Security Policies for Tenant Systems, Security Zones for Tenant Systems, and Security Log for Tenant Systems.]

  - Screen options for attack detection and prevention

    [See Screen Options for Tenant Systems.]

  - Logical tunnel interfaces and GRE tunnels

    [See Flow for Tenant Systems.]

- **UTM support for logical systems (SRX Series)**—Starting in Junos OS Release 18.3R1, unified threat management (UTM) is supported on logical systems. Use the **set security utm default-configuration** command to create a default UTM profile at the master logical system level. You can configure policies, profiles, and custom objects for UTM for each logical system. For a user logical system, parameters such as **mime-whitelist** and **url-whitelist** in an antivirus profile and **address-blacklist** and **address-whitelist** in an antispam profile can be configured at the following hierarchy levels, respectively:

- **[edit security utm feature-profile anti-virus sophos-engine profile]**
- **[edit security utm feature-profile anti-spam sbl profile]**

[See Unified Threat Management Overview.]

- **User firewall support in logical systems (SRX Series)**—Starting in Junos OS Release 18.3R1, user logical systems share user firewall authentication entries such as authentication entry timeout and invalid authentication entry timeout attributes with the master logical system.

  The support for authentication sources is extended to local authentication, Active Directory authentication, and firewall authentication, in addition to the existing supported authentication sources such as Juniper Identity Management Service (JIMS) and Clear Pass authentication.

  [See Overview of Integrated User Firewall.]

NAT

- **NAT configuration check on egress interfaces after reroute (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, support for retaining an existing session with Network Address Translation (NAT) rule is available when there is a change in egress interface because of rerouting.

  If the new egress interface and the previous egress interface are in the same security zone and there is no change in the matched NAT rule or if no rule is applied before and after rerouting, the session is retained with the existing NAT rule. Otherwise, the session expires and new session is created after retransmit or subsequent traffic is received.

  [See Understanding NAT Configuration Check on Egress Interfaces after Reroute.]

- **Session persistence after NAT configuration change (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, SRX Series devices support Network Address Translation (NAT) session persistence. With NAT session persistence enabled on your device, if there are any changes in the NAT configuration, then the device retains the existing NAT sessions instead of clearing them.

  NAT session persistence is supported only for source NAT in the following scenarios:

  - **Source pool**—Change in an address range in a Port Address Translation (PAT) pool.
  - **Source rule**—Change in match conditions for the address book, application, destination IP address, destination port, source IP address, and destination port fields.

  [See Understanding NAT Session Persistence.]

Platform and Infrastructure

- **Juniper Sky ATP Added Platform Support**—Junos OS Release 18.3R1 adds support for SRX300 and SRX320 devices with Juniper Sky ATP.

  [See Juniper Sky ATP Supported Platforms Guide.]

Routing Protocols

- **Support to disable graceful restart helper mode during an interface failure (SRX Series)**—Starting in Junos OS Release 18.3R1, you can prevent SRX Series devices from entering the graceful restart helper mode when the device is configured with BFD with a single-hop external BGP (EBGP).

  To disable the graceful restart helper mode capability, include the **dont-help-shared-fate-bfd-down** statement at the **[edit protocols bgp graceful-restart]** hierarchy level. When the helper mode is not enabled, data traffic continues to be forwarded to an alternate path even if there is an interface failure.

  [See dont-help-shared-fate-bfd-down.]

UTM

- **Explicit proxy support for Enhanced Web Filtering and Sophos antivirus (SRX Series and vSRX)**—Starting in Junos OS Release 18.3R1, SRX Series devices support the use of an explicit proxy for the cloud-based connectivity for Enhanced Web Filtering (EWF) and Sophos antivirus (SAV). It hides the identity of the source device and establishes a connection with the destination device.

  To use the explicit proxy, create one or more proxy profiles and refer to those profiles:

  - In EWF, to establish connection with the Websense Threatseeker Cloud (TSC) server and dynamically load new EWF categories without any software upgrade.
  - In SAV, to connect to the pattern update server using the proxy host IP address.

  [See Understanding Explicit Proxy.]

SEE ALSO

Changes in Behavior and Syntax | 295

Known Behavior | 297

Known Issues | 300

Resolved Issues | 305

Documentation Updates | 315

Migration, Upgrade, and Downgrade Instructions | 315

Product Compatibility | 316

Changes in Behavior and Syntax

IN THIS SECTION

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 18.3R2 for the SRX Series.

Authentication and Access Control

- **Enhanced output for show security firewall-authentication jims statistics (SRX Series)**—Starting in Junos OS Release 18.3R2, the output for **show security firewall-authentication jims statistics** operational command is enhanced to display the statistics of both primary and secondary JIMS server. For example, **show security firewall-authentication jims statistics** operational command displays the following sample output:

```
root@user> show security firewall-authentication jims statistics
```

```
Primary server:
  Push success counter: 1
  Push failure counter: 0

Secondary server:
  Push success counter: 1
  Push failure counter: 0
```

[See show security firewall-authentication jims statistics.]

Chassis Clustering

- **MACsec on Chassis cluster (SRX4600)**—Starting in Junos OS Release 18.3R2, any new MACsec chassis cluster port configurations or modifications to existing MACsec chassis cluster port configurations will require the chassis cluster to be disabled and displays a warning message **Modifying cluster control port CA will break chassis cluster**. Once disabled, you can apply the preceding configurations and enable the chassis cluster.

  [See Configuration Considerations When Configuring MACsec on Chassis Cluster Setup.]

Network Management and Monitoring

- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (SRX Series)**—Starting in Junos OS Release 18.3R2, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, if the server reply after a successful operation includes both an **<ok/>** element and one or more **<rpc-error>** elements with a severity level of warning, the warnings are omitted. In earlier releases, or when the **rfc-compliant** statement is not configured, the NETCONF server might issue an RPC reply that includes both an **<rpc-error>** element with a severity level of warning and an **<ok/>** element.

Platform and Infrastructure

- **Chassis cluster with SPC card (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.3R2, when a SPC is the control plane as well as hosting the control port, this creates a single point of failure. If the SPC goes down on the primary node, the node is automatically rebooted to avoid split brain.

  Connecting SRX Series Devices to Create a Chassis Cluster

VPN

- **Encryption algorithm (SRX Series)**—Starting in Junos OS Release 18.3R2, when AES-GCM 128-bit or AES-GCM 256-bit encryption algorithms are configured in the IPsec proposal, it is not mandatory to configure AES-GCM encryption algorithm in the corresponding IKE proposal.

  [See IPsec VPN Configuration Overview and encryption-algorithm (Security IKE).]

- **Encryption algorithm support for high availability**—Starting in Junos OS Release 18.3R2, on SRX5000 Series devices, you can configure the **aes-128-cbc** option at **set security ipsec internal security-association manual encryption algorithm**. you configure this option for encrypting the high availability link.

  [See internal (Security IPsec).]

- **Certificate revocation list (SRX Series)**—Local certificates are being validated against the certificate revocation list (CRL) even when the CRL check is disabled. Starting in Junos OS Release 18.3R2, this can

be stopped by disabling the CRL check through the Public Key Infrastructure (PKI) configuration. When the CRL check is disabled, PKI will not validate the local certificate against the CRL.

[See revocation-check (Security PKI) and Understanding Online Certificate Status Protocol and Certificate Revocation Lists.]

SEE ALSO

Known Behavior

IN THIS SECTION

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.3R2 for the SRX Series.

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

Application Firewall

- SRX1500 with Application Firewall (AppFW) configured the expected HTTP CPS is 60,000 which is 14 percent drop (expected is 70,000). PR1339131

Chassis Clustering

- On all SRX branch device, if you enable **ip monitoring** for redundancy groups, the feature might not work properly on the secondary node if the reth interface has more than one physical interfaces configured. This is because the backup node will send traffic using the mac address of the lowest port in the bundle. If the reply does not come back on the same physical port, then the internal switch will drop it. PR1344173

Flow-based and Packet-based Processing

-
  - When user configures an interface to a zone under a tenant, interfaces which are rent by other tenant are listed with question mark.
  - When user configures an interface to a zone under root system, interfaces which are rent by other tenant are listed with question mark.

  PR1370255

Interfaces and Chassis

- On SRX4600 devices, USB disk is not made available to Junos. However, the USB disk is available for Host OS (Linux) with full access. USB is still used in the booting process (install and recovery functions). PR1283618

J-Web

- On SRX Series devices, DHCP relay configuration under **Configure > Services > DHCP > DHCP Relay** page is removed from J-Web in Junos OS Release 15.1X49-D60. The same DHCP relay can be configured using the CLI. PR1205911
- On SRX Series devices, DHCP client bindings under Monitor is removed. The same bindings can be seen in CLI using the **show dhcp client binding** command. PR1205915
- PPPoE interface pp0 will not be manageable through J-Web's **Interfaces->Port page**. PR1316328
- When there are no SPC2 card in the SRX5K device, the **Configure > Multi tenancy > Logical systems and Resource Profile** pages will not populate the resource profiles in creation. PR1362106

Unified Threat Management (UTM)

- UTM feature profile works on logical systems level. If mail notify for UTM content filtering is configured on logical systems feature profile level, DUT could not send the mail to the specified mail server due to system SMTP server does not support logical systems. The SMTP server is just accessible for root system feature profile. PR1364783
- From 18.3 release onwards, we started supporting category in APBR module and based on destination IP address, category classification will occur and APBR action will be taken place. UTM web filtering will provide an information about category to APBR module for the matched/received destination IP address. Currently, there is a limitation from web filtering stating that category classification will not be accurate for pure IP address and leads to non-APBR route. PR1365931
- To make APBR custom category to work, we need to create one local utm profile. As a workaround, to create one local utm profile use **set security utm feature-profile web-filtering juniper-local profile h1 category custom action permit** command. PR1366528

User Firewall

- User firewall authentication entries may mismatch when frequently execute command **request services user-identification authentication-table delete authentication-source**. PR1366767

User Interface and Configuration

- In few SRX setups, committing a configuration with a considerable number of logical system configuration can take a little more time than usual. The reason can be taking backup of previous configurations might take a little longer to finish. PR1339862

SEE ALSO

New and Changed Features | **287**

Changes in Behavior and Syntax | **295**

Known Issues | **300**

Resolved Issues | **305**

Documentation Updates | **315**

Migration, Upgrade, and Downgrade Instructions | **315**

Product Compatibility | **316**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.3R2 for SRX Series devices.

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

- After a radius service is restored, SRX device do not send a **RADIUS REQUEST** message to that radius server, which causes authentication requests to time out. PR1366002

Chassis Clustering

- On SRX550 and SRX550M platfroms, SFP-T does not sometimes go up after chassis reboot PR1347874

- On SRX5000, SRX4000, SRX1500 and vSRX platforms with GTP inspection enabled, when updating/deleting a GTP tunnel with GTP-U traffic hitting the tunnel simultaneously, in rare cases, the internal operation of tunnel flag may last for a long time (21 seconds), then the watchdog is triggered causing flowd process to stop. PR1404317
- On all SRX platforms with GPRS tunneling protocol version 2 (GTPv2) traffic logging configuration, the device might be potentially overwritten with an incorrect buffer address if the detailed logging is configured under GTP profile. As a result, it might reboot and cause an outage of the traffic. PR1413718

Flow-Based and Packet-Based Processing

- SRX1500 fan speed often goes up and down if the environment temperature up to 63 degrees celsius. PR1335523
- SRX550M is configured with Layer 2 switching mode and it has an interface configured from an installed GPIM. The interface mode was set to use access mode. VLAN packets may be stripped from the outgoing packet. PR1343394
- SRX1500 devices may encounter a loss in reading/writing access to SSD drive due to an incorrect calculation error during read/write operations with SSD firmware version 560ABBF0. PR1345275
- In a very rare condition, SRX can have a flowd core files when ALG (example- SIP ALG) traffic is crossing firewall and ALG (example- SIP ALG) is enabled. 15.1X49-D140 and above code have a solution for it. PR1352416
- **kern.maxfiles** limit exceeded causes ssh and telnet to the SRX failure. PR1357076
- Application identification classification logic has been improved for NetBIOS and RPC. PR1357093
- On the SRX1500 platforms, the system does not get reset by a watchdog when the CPU freezes. PR1361843
- On all SRX Series platforms, when the flow traceoptions with the packet filter are enabled, the traces of other sessions that are not configured in the packet filter might be captured in the logs. However, when the packet filters are removed, the traces are got dumped into the log file for some time less than 30 seconds. PR1367124
- The reject code for firewall filter is incomplete according to RFC4443 3.1.PR1371115
- Support for intelligent CLI based auto complete was added to secure wire in this release. PR1372825
- With stress TCP traffics, some invalid sessions will timeout over 48 hours. PR1383139
- SRX300 line devices default configuration changed. PR1393683
- On all SRX5000 platforms, when the cluster only has a single SPC card in each node, if the SPC2/SPC3 card goes offline in the primary node, a split brain might occurs. This could cause traffic loss. Reboot both nodes can recover this issue. PR1403872
- On SRX platforms, if knob **enable-session-cache** is configured under the SSL termination scenario, the flowd process might stop. This issue might cause traffic loss. PR1407330

- On all SRX platforms, in chassis cluster with Z mode traffic and local (non-reth) interfaces are configured, when using ECMP routing between multiple interfaces residing on both node0 and node1, if a session is initiated through one node and the return traffic comes in through the other node, packets may get dropped due to reroute failed. PR1410233
- With PMI ON, IPsec encrypted statistics on the RE CLI **show security ipsec statistics** is not working anymore for fragment packets. PR1411486

Forwarding and Sampling

- SRRD process acts as a server for all J-Flow clients. The J-Flow clients can be either PFEs or PICs performing J-Flow. The maximum number of J-Flow clients were previously 32 and it has been increased to 64 in this release. PR1261783

General Routing

- GUMEM errors for the same address might continually be logged if a parity error occurs in a locked location in GUMEM. Since GUMEM utilizes ECC memory, any error is self correcting and has no impact on the operation of the router. In a rare case, such a parity error might appear repeatedly at a specific location. As a workaround, the error can be cleared by rebooting the FPC. PR1200503

J-Web

- CLI Terminal will not be working in Java version 1.8 due to security restriction in running applet. PR1341956
- On SRX platform, root password configured at first J-Web access (**Skip to J-Web** feature) does not work if password length is shorter than 8 characters. PR1371353

Network Address Translation (NAT)

- If UTM web filter is configured and an application is configured into source/destination NAT rule, once this application is deleted or modified, the nsd process might stop. This issue might cause web filter does not work. PR1406248

Network Management and Monitoring

- On all Junos platforms, **etherStatsTable** should display data for IFDs only. However, data from parent IFD (refers to a physical device, see KB2820, https://kb.juniper.net/KB2820) was populated to IFLs (refers to a logical device). To fix this issue, new hidden CLI **set snmp customization ether-stats-ifd-only** has been introduced. When this command is set, snmp walk on **etherStatsTable** will display IFD stats only. PR1335808

Platform and Infrastructure

- On SRX5400, SRX5600, and SRX5600 devices, when the control link is down, the secondary node becomes ineligible and then goes to disabled state. But the FPCs restart continuously after going to disabled state when the FPCs should remain offline until rebooted. PR1170024
- On SRX5600 and SRX5800 models in chassis cluster, when a second routing engine is installed to enable dual control links, the **show chassis hardware** operational command may show the same serial number for both the second routing engines on both the nodes. PR1321502
- On SRX5000 platforms (include SRX5400, SRX5600,SRX5800), EM interface is an internal interface. If EM interface is down that leads to the control link being lost. SRX cluster will be in an abnormal status. PR1342362
- ISSU upgrade from release 15.1X49-D125 to release 17.4X1 might cause multiple flowd process file generates on SRX cluster. PR1363314
- In chassis cluster redundancy group failover scenario, on SRX5600 and 5800 platforms, if the failover is caused by interface monitoring failure, the failover on PFE side (that is data plane) might be slow (example-impact on BFD session up to several seconds). This issue might result in protocol and traffic outage. PR1385521
- On SRX high end platforms, when SPU (Services Processing Unit) VM core happens on one node, this triggers bad kernel state on this node and complete device outage could be seen, which means all IGP and BGP adjacencies would be affected. The reason is that the SPU VM core causes primary PE to dump live VM core, which blocks jsrpd from committing RG (Redundancy Group) state updates to kernel to

set PFE to primary state. And if no PFE is in the primary state traffic would be lost because the original primary SPU is reset in the process of booting up. It is a very rare timing issue. PR1417252

- On **max-vlan-id** that is 1024 particularly for reth interface. when we try to configure more than 1024 we will have the error like **'unit 1026' Too many VLAN-IDs on interface error: configuration check-out failed**. PR1420685

Routing Policy and Firewall Filters

- On all SRX platforms, dns-name entries in policies might not be resolved if the routing instance is configured under a system name server. PR1347006

System Logs

- On SRX5000 series devices between IOC cards and RE. It will prevent repeating of following log message **L2ALM Trying peer/master connection, status 26**. PR1317011

Unified Threat Management (UTM)

- From 18.4 and above release, UTM log will include source and destination zone information. PR1326271

Upgrade and Downgrade

- The issue affects all SRX platforms if doing an ISSU upgrade. The reth interface might flap and cause traffic loss. PR1381475

User Interface and Configuration

- The option of **show command** under tenant context is less than the one under logical systems context. PR1360120

VPNs

- IPSec uses ESP as the default protocol in ipsec proposal, if the protocol is not explicitly configured by the administrator. PR1061838
- When SRX as an initiator behind the NAT, disabling NAT on the middle router causes an immediate new negotiation failure due to an attempt with port 4500. The next attempt will succeed by using port 500. Disabling NAT and bringing down all the existing tunnels and re-establishing the tunnels with port 500 is the expected behavior solution. PR1273213

- VPN tunnels flap after adding or deleting a configuration group in **edit private** mode on a clustered setup. PR1400712
- VPN does not recover on the high-end standalone SRX when CLI operation **restart ipsec-key-management** is done. PR1390831
- On SRX5400, SRX5600, SRX5800 devices with SPC3, idle IPsec VPN tunnels without traffic and with ongoing DPD probes will be affected during the RG0 failover window. IPsec VPN daemon in the new primary routing-engine may not be initialized on-time to respond to the DPD probes. PR1405515

SEE ALSO

New and Changed Features | **287**

Changes in Behavior and Syntax | **295**

Known Behavior | **297**

Resolved Issues | **305**

Documentation Updates | **315**

Migration, Upgrade, and Downgrade Instructions | **315**

Product Compatibility | **316**

Resolved Issues

This section lists the issues fixed in hardware and software in Junos OS Release 18.3R2 for SRX Series devices.

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

Resolved Issues: 18.3R2

Application Layer Gateways (ALGs)

- On all SRX platforms, SIP/FTP ALG does not work when SIP traffic with source NAT goes through SRX device. PR1398377

Application Security

- Fail to match permit rule in application firewall ruleset. PR1404161

Application Layer Gateways (ALGs)

- DNS requests with EDNS options might be dropped by DNS ALG. PR1379433

- The SUN-RPC data traffic might be dropped after interface related configuration is changed. PR1387895
- H323 voice packets might be dropped on SRX devices. PR1400630

Chassis Clustering

- The half duplex mode do not support on SRX340 and SRX345 PR1149904
- Multiple flowd process files are seen on node1 after an RG0 failover. PR1372761
- Traffic loss occurs when the primary node is rebooting. PR1372862
- The packet might be dropped in an SRX chassis cluster environment if the sampling or packet capture is configured. PR1379734
- The flowd might stop if doing an ISSU upgrade. PR1386522
- VDSL is not stable if there are sudden noises after configuring VDSL SOS feature. PR1387133
- If using SRX cluster and configuring 4 100G interfaces on PIC 0, all the 4 interfaces might be down. PR1387701
- ISSU status with error from 18.2R1-S1/18.2R1-S2 to 18.2R1-S3. PR1387947
- The cluster IDs larger than 10 will cause FPCs to remain in offline on SRX4600 chassis cluster. PR1390202
- The MACsec on a physical port may not initialize properly when a new node is joined to chassis cluster. PR1396020
- Traffic with domain name address might fail for 3-5 minutes after RG0 failover on SRX platforms. PR1401925

Command-Line Interface (CLI)

- Display issue in **show usp memory segment shm data module** and **show jsf shm module vty fwdd** commands on branch SRX. PR1387711

Flow-based and Packet-based Processing

- Security Logs for unified policies have been improved to correctly reflect the reason for a denied or rejected session. PR1338310
- Crash happens when the output interface is configured in X2 mirror filter configuration is down. PR1357347
- Control traffic loss may be seen on SRX4600 platform. PR1357591
- Application identification support for HTTP, SMTPS, POP3S, and IMAPS applications. PR1365810
- SRX1500 continues alarm on FAN **Fan Tray 0 Fan 0 Spinning Degraded**. PR1367334
- Observing 40 percent drop with respect to basic FW UDP IMIX throughput(expected is 20 percent) PR1373019
- PIM register message might be dropped on SRX Series devices. PR1378295

- The pkid process might stop after RG0 failover. PR1379348
- On SRX1500, activity LED (right LED) for 1G/10G port is not on although traffic is passing through that interface. PR1380928
- Improper message is thrown when the data path debug capture is stopped. PR1381703
- SRX5600 HA ICAP redirect status flapping on few SPU PICs. PR1382376
- The flowd/srxpfe process might crash when SSL proxy is used. PR1383655
- Large file downloads slow down for many seconds. PR1386122
- Traffic might be processed by the VRRP backup when multiple VRRP groups are configured. PR1386292
- Traffic might be stopped after session created on SRX4600 platform. PR1388735
- The SRX does not send messages **frag needed and DF set** back to the source host during path MTU discovery. PR1389428
- Future group membership updates are not recognized by IUFW after a users sAMAccountName is changed while its DN remained the same. PR1394049
- Packet loss might occur on unrelated traffic when AppQoS rate limiter is applied on SRX4600 and SRX5000 Series platform using SPC3. PR1394085
- These messages are seen **/kernel: tcp_timer_keep:Local(0x80000004:54652) Foreign(0x80000004:33160)**. PR1396584
- Request to unhide **dropped-illegal-packet** and **dropped-icmp-packet** configuration options. PR1394720
- Switching interface mode between **family ethernet-switching** and **family inet/inet6** might cause traffic loss. PR1394850
- SRX connection to JIMS keeps flapping causes failover to secondary JIMS. PR1398140
- On SRX4600 and SRX5000 Series devices, BGP packets might be dropped under high CPU usage. PR1398407
- VLAN push might not work on SRX1500. PR1398877
- Increase DAG feed scale number to 256 from 63. PR1399314
- Unable to access to SRX platforms if messages **kern.maxfiles limit exceeded by uid 65534, please see tuning(7)** are seen. PR1402242
- Downloads may stall and/or completely fail when utilizing services that are reliant on TCP proxy. PR1403412
- Transit UDP 500/4500 traffic might not pass across SRX5000 series devices when using SPC3/SPC2. PR1403517
- ISSU failed from 18.3R1.9 to 18.4R1.4. PR1405556
- The flowd process stops and all cards are brought off. PR1406210
- The RG1 failover does not happen immediately when the SPC3 card crashes. PR1407064

- IDP signature update fails at RG0 primary node. PR1407603
- Memory leak if AAMW is enabled. PR1409606
- Session capacity of SRX340 is not match SRX345. PR1410801
- Any traffic originated from the device itself might be dropped in the IPsec tunnel. PR1414509
- Command **show security firewall-authentication jims statistics** will output statistics of both primary jims server and secondary jims server. PR1415987
- Traffic logging shows service-name **junos-dhcp-server** for UDP destination port 68. PR1417423

General Routing

- High jsd or na-grpcd CPU usage might be seen even JET or JTI is not used. PR1398398
- The authd might stop when issuing **show network-access requests pending** command during the authd restarting. PR1401249

Intrusion Detection and Prevention (IDP)

- Unable to deploy IDP due to the IDP configuration cannot be committed. PR1374079
- When utilizing unified policies with IDP, under certain circumstances IDP would not inspect arbitrary sessions, marking them as **Not Interested** within **show security idp counters flow**. PR1385094
- Performance drops are seen in SRX345/SRX340 platforms for IDP C2S policy. PR1395592
- Unable to configure dynamic-attack-group. PR1418754

Interfaces and Routing

- SRX1500 IPv4 multicast packets might not be broadcasted from the IRB interface. PR1385934
- SRX4600 10G Interface optics diagnostic access issue. PR1395806
- 40G/100G ports may take a long time (about 30 seconds) to link up on SRX4600 platform. PR1397210
- SRX device can not obtain IPv6 address through DHCPv6 when using PPPOE interface with logical-unit-number greater than zero. PR1402066

Installation and Upgrade

- 18.3R1 cannot be installed through TFTP in boot loader on SRX 300 line platforms. PR1390858

J-Web

- In the J-Web Dashboard, the **Security Resources** widget did not display absolute values. PR1372826
- Excluded addresses within J-Web security policy editor were not sufficiently differentiated versus normal addresses. They are now highlighted red for ease of identification. PR1376112
- In this release, J-Web now supports defining SSL-Proxy and redirect (block page) profiles when a policy contains dynamic applications. PR1376117
- Chassis image did not show from J-Web dashboard PR1382219

- J-Web page do not load after login with logical-system specific user. PR1396879
- The next-hop IP address is not displayed in the routing table in the J-Web. PR1398650
- Special character used in the pre-shared-key is removed silently after a commit operation on J-Web. PR1399363
- Configuring using the CLI editor in the J-Web generates an mgd core file. PR1404946
- The httpd-gk process stops leading to dynamic VPN failures and high RE CPU utilization 100 percent. PR1414642

Layer 2 Ethernet Services

- DHCPv6 clients might fail to get addresses on SRX platforms. PR1392723

Logical Systems and Tenant Systems

- Logical system license fail to bind to the tenant/logical systems after rebooting the device. PR1380144
- Logical system license. PR1384659
- Logical system configuration installed failed on node 1 after ISSU from 18.2R1.9 to 18.3R1.8. PR1388336

Network Address Translation (NAT)

- The SRX might send the **noSuchInstance** value to SNMP server in get response during commit. PR1357840
- NAT64 and traceroute do not work correctly on an SRX. PR1376890
- SRX-SPC3 mix mode NAT SPC3 core at **../sysdeps/unix/sysv/linux/raise.c:55**. PR1403583

Platform and Infrastructure

- High httpd utilization after reboot failover. PR1352133
- Many chassis commands missing. PR1363645
- Packet capture feature does not work after removing the sampling configuration. PR1370779
- IP monitoring failure resulting in multiple interfaces disappearing from forwarding table. PR1371500
- Some error messages could be seen when running **show interface extensive** command from CLI or Junos Space. PR1380439
- Traffic loss seen in Layer 2 VPN with GRE tunnel. PR1381740
- Junos upgrade might fail with validate option after the **/cf/var/sw** directory is accidentally deleted. PR1384319
- Login class with allowed days and specific access-start/access-end does not work as expected. PR1389633
- GW lcores and srxpfe cores at **../src/pfe/usp/rt/applications/ipsec/ipsec_rt_forge_util.c:59** when loading 18.4 image. PR1392580

- The flowd process stops if it goes into a dead loop. PR1403276
- RE CPU utilization is high and eventd is consuming a lot of resources. PR1418444

Routing Policy and Firewall Filters

- The output of **show security policies detail** has been modified to improve readability, particularly for unified policies. PR1338307
- The timeout value of junos-http is improper. PR1371041
- When SSL-Forward-Proxy is configured in a unified policy along with the action of Reject+Redirect, a block page was not presented to the user for HTTPS sites PR1375823
- **show security flow session** now fully supports the dynamic-application construct. PR1387449
- The nsd process stops and generates a core file. PR1388719

Services Applications

- Flowd process stops in icap_redirect_release_profile_server at **../../../../../../../src/pfe/usp/rt/applications/icap-redirect/icap_redirect_server.c:1513**. PR1389600
- SRX5600 HA SPC2 ICAP redirect object's are in use even after clearing TCP sessions. PR1390835

Unified Threat Management (UTM)

- EWF server status shows **UP** when 443 is specified as server port. PR1383695
- Whitelist/Blacklist does not work for HTTPS traffic going through Web proxy. PR1401996
- UTM Web filtering status shows down when using Hostname [routing-instance synchronization failure]. PR1421398

VPNs

- The kmd process might stop when configuring IPsec VPN and BGP on SRX1500 platform. PR1336235
- Dot usage in CA profile name causes issues when the pkid process is restarted. PR1351727
- SPC3 ike sa detail output is not showing proper traffic statistics. PR1371638
- In a rare situation, VPN tunnels may not be configured successfully and the VPN tunnels will not come up. PR1376134
- Packet loss was seen in IPsec Z-mode scenario. PR1377266
- The kmd daemon might stop and cause VPN traffic outage after executing **show security ipsec next-hop-tunnels**. PR1381868
- Adding/deleting site-to-site manual-nhtb VPN tunnels to an existing st0 unit will cause existing manual-NHTB VPN tunnels under the same st0 unit to flap. PR1382694
- After repeatedly HA failover, the flowd process might stop if IPSec VPN is configured. PR1386229

- On SRX5400, SRX5600, SRX5800 devices with SPC3, **show security ike security-association detail** command does not display local IKE-ID field correctly. PR1388979
- A few VPN tunnels do not forward traffic after RG1 failover. PR1394427
- The kmd process might stop when SNMP polls for the IKE SA. PR1397897
- Syslog is not generated when ike gateway rejects duplicate IKE ID connection. PR1404985
- Not all the tunnels are deleted when authentication algorithm in ipsec proposal is changed. PR1406020
- Multiple flowd process files are observed with IPsec acceleration with fragmentation traffic. PR1407910
- Traffic drops on peer due to bad SPI after first re authentication. PR1412316

Resolved Issues: 18.3R1

Application Layer Gateways (ALGs)

- When using IPsec ALGs, the IPsec tunnel payload is dropped after IKE/IPsec tunnel reestablishment because of session conflict. PR1372232
- The status of SIP ALG is disabled and the original SIP active sessions are affected, when SIP active sessions are created with standard port 5060. PR1373420

Application Layer Gateways (ALGs)

- On SRX5800 devices when IPsec ALG is used, the IPsec tunnel payload is dropped after IKE/IPsec tunnel reestablishment because of session conflict. PR1372232
- When the status of SIP ALG is changed to disabled, the SIP active session is affected. PR1373420
- DNS requests with additional EDNS records might be dropped by the DNS ALG. PR1379433

Class of Service (CoS)

- Packets go out of order on SPC2 cards when IOC1 or FIOC cards are used. PR1339551

Flow-Based and Packet-Based Processing

- Using SSH to connect to the loopback interface of the SRX Series device does not work properly when AppTrack is configured. PR1343736
- SNMP MIB walk provides wrong data counters for total current flow sessions. PR1344352
- File download stops over a period of time when TCP proxy is activated through antivirus or Juniper Sky ATP. PR1349351
- When the routing instance is configured, the UTM Anti-Spam:DUT process does not send the DNS query. PR1352906
- IPsec VPN traffic might drop when passing through the SRX Series device after an IKE rekey. PR1353779

- IPv6 backup sessions might hang and cannot be cleared after data-plane redundancy groups fail over. PR1354448
- The PIM register message might stop from the source first-hop router. PR1356241
- On the SRX5000 line of devices, when the IPsec performance acceleration feature is enabled, packets going in to or out of a VPN tunnel are dropped. PR1357616
- On the secondary control plane, a multicast session leak is observed when the PIM is registered. PR1360373

Interfaces and Chassis

- On SRX4600 device, the virtual IP address of the VRRP might not respond to host-inbound traffic. PR1371516

Intrusion Detection and Prevention (IDP)

- Unable to load IDP policy because of less available heap memory. PR1347821
- IDP signature update fails on secondary node. PR1358489

J-Web

- The Dynamic-Application configuration page does not display application signatures properly when you search using the category filter. PR1344165
- In J-Web you cannot delete dynamic VPN user configuration. PR1348705
- When J-Web fails to get resource information, the Routing Engine CPU usage shows 100% resource utilization on the J-Web dashboard. PR1351416
- When you use Internet Explorer version 11, the security policies search button in J-Web does not work. PR1352910
- J-Web setup wizard does not propagate DHCP attributes from ISP to LAN. PR1370700

Layer 2 Features

- The dcpfe and fxpc processes might stop on Packet Forwarding Engines with low memory. PR1362332

Layer 2 Ethernet Services

- The subnet mask is not sent as the reply to a DHCPINFORM message. PR1357291

Network Management and Monitoring

- With user firewall enabled and RG0 failover is being performed, eventd process core files are generated. PR1366120

Platform and Infrastructure

- VPN is not stable when you perform commits with apply-groups. PR1242757
- The **show chassis environment pem** and **show chassis power** commands show incorrect input voltage. PR1323256
- On SRX Series devices, the **No Port is enabled for FPC# on node0** log is generated every 5 seconds. PR1335486
- On the SRX5000 line of devices, frequent logs are seen when the IOC has the same identifier as the SPC PIC. PR1357913
- On SRX4100 devices, the interface shows up as half-duplex. PR1358066
- SCP configuration backup fails even though **/var/etc/ssh_known_hosts** has the correct fingerprint. PR1359424

Routing Policy and Firewall Filters

- The flowd process stops after a large number of custom applications are configured. PR1347822
- On SRX Series devices, the nsd process might stop on the Packet Forwarding Engine with large-scale security policy configuration. PR1354576
- Dynamic application autocomplete support is not functional within the CLI for the **show security match-policies** command. PR1363908
- The timeout value of **junos-http** is incorrect. PR1371041
- When a policy references dynamic addresses in the destination-address field and the destination IP address of the traffic is within this dynamic-address pool, the policy does not match this traffic. The issue occurs only for destination address and not for the source address. PR1372921

Routing Protocols

- When BGP traceoptions are configured and enabled, the traces specific to messages sent to the BGP peer (BGP SEND traces) are not logged, but the traces specific to received messages (BGP RECV traces) are logged correctly. PR1318830
- The ppmd process might stop during ISSU. PR1347277

- On SRX1500 devices, dedicated BFD does not work. PR1347662

Unified Threat Management (UTM)

- The default action of Web filtering does not work as expected. PR1365389

VLAN Infrastructure

- On SRX Series devices in transparent mode, the flowd process might stop when matching the destination MAC address. PR1355381

VPNs

- IPsec traffic statistics counters return 32-bit values, which is too fast and might overflow. PR1301688
- The kmd process might stop if multiple IKE gateways use the same IKE policy. PR1337903
- On the SRX5000 line of devices in a chassis cluster, control link encryption does not work. PR1347380
- After a chassis cluster failover, all IPsec tunnels that are in active state are shown as inactive. PR1348767
- On SRX Series devices, the policy-based IPsec VPN does not forward traffic properly when ingress and egress interfaces are in a virtual router. PR1350123
- On SRX Series devices in a chassis cluster, configuration commit might succeed even when the external logical interface configuration (reth) associated with the IKE VPN gateway configuration is deleted. This might lead to configuration load failure during the next device bootup. PR1352559
- S2S tunnels are not redistributed after IKE and IPsec are reactivated in the configuration. PR1354440
- On SRX5000 line of devices, during the migration from site-to-site VPN to AutoVPN configuration, loss of traffic for some sessions might be observed. PR1362317

SEE ALSO

New and Changed Features | **287**

Changes in Behavior and Syntax | **295**

Known Behavior | **297**

Known Issues | **300**

Documentation Updates | **315**

Migration, Upgrade, and Downgrade Instructions | **315**

Product Compatibility | **316**

Documentation Updates

There are no errata or changes in Junos OS Release 18.3R2 for the SRX Series documentation.

SEE ALSO

- New and Changed Features | 287
- Changes in Behavior and Syntax | 295
- Known Behavior | 297
- Known Issues | 300
- Resolved Issues | 305
- Migration, Upgrade, and Downgrade Instructions | 315
- Product Compatibility | 316

Migration, Upgrade, and Downgrade Instructions

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life 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 15.1X49, 17.3, 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from one Junos OS Release to the next release or one release after the next release. For example you can upgrade from Junos OS Release 15.1X49 to Release 17.3 or 17.4, Junos OS Release 17.4 to Release 18.1 or 18.2, and from Junos OS Release 18.1 to Release 18.2 or 18.3 and so on.

You cannot upgrade directly from a non-EEOL release to a release that is more than three releases ahead or behind. To upgrade or downgrade from a non-EEOL release to a release more than three releases before

or after, first upgrade to the next EEOL release and then upgrade or downgrade from that EEOL release to your target release.

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

For information about software installation and upgrade, see the Installation and Upgrade Guide for Security Devices.

For information about ISSU, see the Chassis Cluster User Guide for Security Devices.

SEE ALSO

- New and Changed Features | 287
- Changes in Behavior and Syntax | 295
- Known Behavior | 297
- Known Issues | 300
- Resolved Issues | 305
- Documentation Updates | 315
- Product Compatibility | 316

Product Compatibility

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

SEE ALSO

- New and Changed Features | 287
- Changes in Behavior and Syntax | 295
- Known Behavior | 297

Upgrading Using ISSU

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.

For additional information about using ISSU on routing and switching devices, see the High Availability User Guide.

For additional information about using ISSU on security devices, see the Chassis Cluster User Guide for SRX Series 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 the Junos OS, see the Juniper Networks Problem Report Search application at https://prsearch.juniper.net.

For regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6 for Juniper Networks products, see the Juniper Networks Compliance Advisor.

To access Software Release Notifications for Junos OS Service Releases, visit our Knowledge Center at https://support.juniper.net/support/. You'll need to log in to your Juniper Account. From the Knowledge Center, search by the specific release number, for example 17.4R1-S2. Use the Software Release Notifications to download software, and learn about known and resolved issues for specific service releases.

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

Documentation Feedback

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

- Online feedback system—Click TechLibrary Feedback, on the lower right of any page on the Juniper Networks TechLibrary site, and do one of the following:

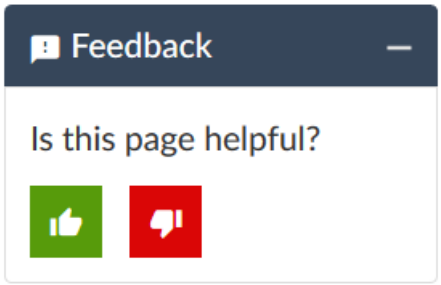

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- E-mail—Send your comments to techpubs-comments@juniper.net. Include the document or topic name, URL or page number, and software version (if applicable).

Requesting Technical Support

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

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

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

Self-Help Online Tools and Resources

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

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

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Opening a Case with JTAC

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

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

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

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

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

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

Revision History

26 August 2021—Revision 7, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 6, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 January 2020—Revision 5, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 October 2019—Revision 4, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 October 2019—Revision 3, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 July 2019—Revision 2, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 May 2019—Revision 1, Junos OS Release 18.3R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 March 2019—Revision 13, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 January 2019—Revision 12, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 January 2019—Revision 11, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 January 2019—Revision 10, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 December 2018—Revision 9, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 December 2018—Revision 8, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 November 2018—Revision 7, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 November 2018—Revision 6, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 November 2018—Revision 5, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 October 2018—Revision 4, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 October 2018—Revision 3, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 October 2018—Revision 2, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 September 2018—Revision 1, Junos OS Release 18.3R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.