Release Notes: Junos® OS Release 18.2R3 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

28 January 2022

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 18.2R3 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.

NOTE: The recommended release for Junos Fusion Data Center is 18.1R2-S2. The subsequent 18.xRx mainline releases (18.2, 18.3, and 18.4) do not support Junos Fusion Data Center.

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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 18.2R3 New and Changed Features

There are no new and changed features in Junos OS Release 18.2R3 for ACX Series Universal Metro Routers.

Release 18.2R2 New and Changed Features

There are no new and changed features in Junos OS Release 18.2R2 for ACX Series Universal Metro Routers.

Release 18.2R1 New and Changed Features

Hardware

- **New ACX5448 Universal Metro Routers**—Starting with Junos OS Release 18.2R1, the ACX5448 Universal Metro Routers are available as Juniper Networks' top-of-rack router solutions for data centers and campus distribution or aggregation environments. The ACX5448 router portfolio consists of high-performance fixed-configuration routers that add higher port densities, additional scalability, and improved latency to the ACX Series. The ACX5448 routers offers a compact 1U model that provides wire-speed packet performance, very low latency, and a rich set of Layer 2 and Layer 3 features. The router has a high-throughput Packet Forwarding Engine, and the performance of the control-plane running on ACX5448 router is enhanced by the 1.9 Ghz six-core Intel CPU with 32 GB of memory and two 100 GB of solid-state drive (SSD) storage.

  The ACX5448 is a 10-Gigabit Ethernet enhanced small form-factor pluggable (SFP+) top-of-rack router with 48 SFP+ ports, and four 100-Gigabit Ethernet QSFP28 ports. Each SFP+ port can operate as a native 10-Gigabit Ethernet port, or as a 1-Gigabit Ethernet port when 1-Gigabit optics are inserted.

The ACX5448 is shipped with redundant fans and redundant power supplies. The router can be ordered with front-to-back airflow (air out or AFO), or back-to-front airflow (air in or AFI), and with AC or DC power supplies.

- **ACX6360 with CFP2-DCO-T-WDM-1 transceiver**—Starting in Junos OS Release 18.2R1, the ACX6360 universal metro router supports the CFP2-DCO-T-WDM-1 transceiver. The following modulation formats are supported on this transceiver:
  - QPSK-100G
  - 8QAM-200G
  - 16QAM-200G

  The ACX6360 router with CFP2-DCO-T-WDM-1 transceiver supports minimum channel spacing of 6.25Ghz.

  [See Understanding the features of ACX6360.]

Authentication Access Control

- **Enhancement to NTP authentication method (ACX500, ACX1100)**— Starting in Junos OS Release 18.2R1, Junos OS supports NTP authentication for both SHA-1 and SHA2-256, in addition to the existing NTP authentication method, MD5. You can now choose from among MD5, SHA-1, and SHA2-256 for synchronizing the clocks of Juniper Network routers, switches, and other security devices on the Internet. Using SHA-1 instead of MD5 improves the security of devices with very little impact to timing, while using SHA2-256 provides an increase in security over SHA-1.

  > **NOTE:** By default, network time synchronization is unauthenticated.

  To implement authentication, use **set authentication-key <*key_number*> type** at the **[edit system ntp]** hierarchy level.
  - To enable SHA-1 authentication, use **set authentication key <*key_number*> type sha1 value <*password*>** at the **[edit system ntp]** hierarchy level.
  - To enable SHA2-256 authentication, use **set authentication key <*key_number*> type sha256 value <*password*>** at the **[edit system ntp]** hierarchy level.

  [See authentication-key and Configuring NTP Authentication Keys.]

Class of Service (CoS)

- **Support for logical interface-based classification and rewrites (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports configuring logical interface-based classification and rewrite rules. ACX5448 router supports fixed, behavior aggregate (IP precedence, DSCP, DSCP IPv6, MPLS EXP, IEEE-802.1p, IEEE-802.1ad (DEI bit)), and multifield classifiers.

  [See Classifiers and Rewrite Rules at the Global, Physical and Logical Interface Levels Overview.]

- **Support for port-based queueing, scheduling, and shaping (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports port-based queueing, scheduling, and shaping. You can configure up to eight queues (virtual output queues) per physical interface (port). Scheduling properties can be applied at both physical as well as logical interface levels. The egress scheduler supports two priority levels (**strict-high** and **low**). Multiple strict-high priority queues and multiple low (default) priority queues can be configured.

  Schedulers and their associated shapers control the traffic bandwidth, jitter (delay variation), and packet loss priority at the egress of the device. By default a port on ACX5448 router gets a dedicated buffer of 100 microseconds and shared buffer from DRAM. Delay buffer controls the latency of the queue during congestion and maximum number of packets that can be held in a queue. Default buffer size per port is 100 microseconds.

  [See Understanding Schedulers Overview, Configuring Shared and Dedicated Buffer Memory Pools, and Hierarchical Class of Service in ACX5000.]

Dynamic Host Configuration Protocol

- **Support for DHCPv4 and DHCPv6 (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports DHCP server, DHCP client, and DHCP relay configuration for IPv4 and IPv6 services. You can enable ACX5448 router to function as DHCP server and configure the DHCP server options on the router. The DHCP server provides an IP address and other configuration information in response to a client request. DHCP relay agent forwards DHCP request and reply packets between a DHCP client and a DHCP server.

  [See Extended DHCP Local Server Overview and Extended DHCP Relay Agent Overview.]

EVPN

- **Support for VPWS with EVPN signaling mechanisms and flexible cross connect (ACX5448)**—Starting with Junos OS Release 18.2R1, the ACX5448 router supports VPWS with EVPN signaling mechanisms and flexible cross connect. The EVPN VPWS provides a framework for delivering the VPWS with EVPN signaling mechanisms. The VPWS with EVPN signaling mechanisms supports single-active or all-active multihoming capabilities and inter-autonomous system (AS) options associated with BGP-signaled VPNs. The EVPN VPWS flexible cross connect addresses the label resource issue. The flexible cross-connect (FXC) service enables interoperability of access router that uses EVPN FXC VLAN-aware and VLAN-unaware FXC services. ACX5448 router do not support pseudowire services in EVPN VPWS flexible cross connect.

  The following limitations apply:

  - Control word is not supported for EVPN VPWS services.
  - When VLAN maps are applied on the ccc-interfaces (UNI) for EVPN VPWS, only the following VLAN map operations are applicable:

```
IFD Encap/      IFL-TYPE         Input-MAP       Output-Map
-------------------------------------------------------------
ethernet-ccc    unit 0;
                TC2             push-push       pop-pop

vlan-ccc
                ST: vlan-id X
                                swap-push        ..
                DT: vlan-tags outer  X  inner Y
                TC1             pop-pop         push-push
                TC4             swap-swap       swap-swap
```

- VLAN map with non-default TPIDs in the VLAN map operation is not supported.
- Aggregated Ethernet interfaces with LAG interface for EVPN VPWS and EVPN VPWS FXC services are not supported. However, for CE multihoming, the CE can have static-AE, and CE can multihome to the ACX5448 PE router (PE in non-AE/LAG).

[See Overview of VPWS with EVPN Signaling Mechanisms.]

General Routing

- **Support for virtualization (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 routers support virtualization. Virtualization enables multiple instances of operating systems, called guests, to run concurrently on the host and share virtualized hardware resources. A guest is a virtual machine (VM) that runs on a hypervisor-based host and shares its resources. A host is a virtualized software whose hypervisor allows multiple guest VMs to run on it concurrently and share its resources. A VM can be an instance of Junos OS or any compatible third-party VM. Each VM runs its own operating system image and applications that can be different from that of another VM running on the same host. ACX5448 router supports only one Junos VM. You can use the following chassis management commands to manage the onboard FRUs:
  - **show chassis hardware**
  - **show chassis temperature-thresholds**
  - **show chassis environment**
  - **show chassis alarms**

  ACX5448 router emulates one FPC with two PICs. One PIC represents the 48x1/10GE ports and other represents the 4x100GE ports. The **show chassis hardware** CLI command shows the FPC and PICs as built-in as shown in the following sample output:

  user@host> **show chassis hardware**

```
Hardware inventory:
Item             Version  Part number  Serial number     Description
Chassis                                DA805             ACX5448
Midplane         REV 13   750-065110   ACNP4346          ACX5448
Routing Engine            BUILTIN      BUILTIN           Routing Engine
RFEB
FPC 0                     BUILTIN      BUILTIN           FPC BUILTIN
  MIC 0                                                  48x1GE/48x10GE
    PIC 0                 BUILTIN      BUILTIN           48x1GE/48x10GE
  MIC 1                                                  24x10/25GE 6x40/100GE
    PIC 1                 BUILTIN      BUILTIN           24x10/25GE 6x40/100GE
```

NOTE: ACX5448 routers do not support **request system software upgrade** and **request system software rollback** commands, instead you must use **request vmhost** CLI commands.

ACX5448 routers do not support:

- Multiple guest VMs
- Redundant Junos VMs
- ISSU
- 10/100 Mbps copper SFPs

[See Routing Engines with VM Host Support and Architecture of Routing Engines with VM Host Support.]

Interfaces and Chassis

- **Resiliency support (ACX6360)**—Starting with Junos OS Release 18.2R1, software resiliency is enabled for ACX6360 routers.
- **Support for 48x1/10GE and 4x100GE Interface Ports (ACX5448)**—ACX5448 router has 48x1/10GE interface ports and 4x100GE interface ports. The 48 ports on ACX5448 router can be configured as 1GE or 10GE modes and these ports are represented by **xe** interface type. The PIC 1 of FPC 0 has 4x100GE ports, where each port can be channelized as 1x100GE, or 1x40GE, or 4x25GE modes and these ports are represented by **et** interface type. By default, the port speed in PIC 1 is 100GE.

  [See Understanding Interfaces on ACX Series Universal Metro Routers.]

Layer 2 Features

- **Support for Layer 2 features (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports the Layer 2 bridging, Q-in-Q tunneling, no-local switching, and Layer 2 protocol tunnel. A bridge domain is a set of logical interfaces that share the same flooding or broadcast characteristics. Layer 2 logical interfaces are created by defining one or more logical units on a physical interface with

encapsulation as **vlan-bridge** and as **ethernet-bridge**. All the member ports of the bridge domain participate in Layer 2 learning and forwarding. These bridging features are used to configure E-LINE, E-LAN and E-TREE services. On ACX5448 router, you can configure bridge domains by using the following methods:

- Bridge domain without a vlan-id number statement
- Bridge domain with the vlan-id value set to none
- Bridge domain with a single vlan-id

The Layer 2 Next Generation mode, also called Enhanced Layer 2 Software (ELS), is supported on ACX5448 router for configuring the Layer 2 features.

If **no-local-switching** is configured in a bridge domain, then traffic cannot flow between CE to CE interfaces. This includes known unicast/multicast, unknown unicast/multicast, and broadcast traffic. However, traffic can flow between CE to PE interfaces and between PE to PE interfaces.

Q-in-Q tunneling allows you to create a Layer 2 Ethernet connection between two customer sites. Providers can segregate different customers' VLAN traffic on a link (for example, if the customers use overlapping VLAN IDs) or bundle different customer VLANs into a single service VLAN. Service providers can use Q-in-Q tunneling to isolate customer traffic within a single site or to enable customer traffic flows across geographic locations.

Layer 2 protocol tunnel can be configured on the customer edge port using mac rewrite configuration. MAC rewrite is supported for the STP, CDP, VTP, LLDP, ELMI, 802.1x, 802.3ah, LACP, MMRP, MVRP protocol packets.

[See Layer 2 Bridge Domains on ACX Series Overview, Q-in-Q Tunneling on ACX Series Overview, and Understanding Layer 2 Next Generation Mode on ACX Series Routers.]

- **Support for Layer 2 services (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports configuring Layer 2 services such as RSTP, MSTP for loop resolutions, and storm control to monitor traffic levels and to drop broadcast, unknown unicast, and multicast (BUM) packets if they exceed the configured limit.

  Storm control is applied on the following traffic types:

  - Layer 2 multicast packets
  - Layer 2 unregistered multicast packets
  - Layer 2 registered multicast packets

  On ACX5448 router, storm control is only applicable at the physical interface level. No event will be logged when a traffic storm hits an ACX5448 router. Also interfaces will not be bound to any default profile. The default action is to drop the packets exceeding the configured bandwidth.

  [See Storm Control on ACX Series Routers Overview.]

- **Support for Layer 2 protection (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports configuring bridge protocol data unit (BPDU) protect, loop protect, and root protect on

spanning-tree instance interface. You can configure BPDU protection on individual interfaces or on all the edge ports of the bridge.

[See Understanding BPDU Protection for Spanning-Tree Instance Interfaces, Understanding Loop Protection for Spanning-Tree Instance Interfaces, and Understanding Root Protection for Spanning-Tree Instance Interfaces in a Layer 2 Switched Network.]

Layer 3 Features

- **Support for Layer 3 features (ACX5448)**—Starting with Junos OS Release 18.2R1, the ACX5448 router uses MPLS as a transport mechanism and they include support for label-switching router (LSR), label edge routers (LERs), and pseudowire services. The protocols such as ECMP, OSPF, ISIS, and BGP are also supported on ACX5448 router.

  [See MPLS Overview.]

Management

- **Support for NETCONF over SSH and custom YANG models (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports NETCONF OVER SSH and custom YANG modules.

  Client applications can access the NETCONF server using the SSH protocol and use the standard SSH authentication mechanism. After authentication, the NETCONF server uses the configured Junos OS login usernames and classes to determine whether a client application is authorized to make each request.

  You can load custom YANG modules on the router to add data models that are not natively supported by Junos OS but can be supported by translation. Doing this enables you to extend the configuration hierarchies and operational commands with data models that are customized for your operations. You can load custom YANG modules by using the **request system yang add** operational command.

  [See Establishing an SSH Connection for a NETCONF Session and YANG Modules Overview.]

MPLS

- **Support for MPLS ping and Bidirectional Forwarding Detection over virtual circuit connection verification (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports MPLS ping and Bidirectional Forwarding Detection over Virtual Circuit Connection Verification. MPLS ping functionality diagnoses the state of label-switched paths (LSPs), where the router sends probe packets into the LSP. Based on how the LSP at the remote endpoint of the connection replies to the probes, you can determine the connectivity of the LSP. Each probe is an echo request sent to the LSP as an MPLS packet with a UDP payload. If the outbound node receives the echo request, it checks the contents of the probe and returns a value in the UDP payload of the response packet. If the Junos OS receives the response packet, it reports a successful ping response.

  Bidirectional Forwarding Detection (BFD) support for virtual circuit connection verification (VCCV) allows you to configure a control channel for a pseudowire, in addition to the corresponding operations and management functions to be used over that control channel. BFD provides a low resource mechanism for the continuous monitoring of the pseudowire data path and for detecting data plane failures, as described in RFC 5885.

You can use the following commands for debugging:

- show bfd session extensive
- show ldp database extensive

[See Pinging LSPs and Configuring BFD for VCCV for Layer 2 Circuits.]

- **Support for MPLS ping and traceroute (ACX5448)**—Starting with Junos OS Release 18.2R1, the ACX5448 router supports MPLS ping and traceroute. MPLS ping and traceroute [RFC-4379] are common tools used to debug connectivity between two PEs for a LSP. The ping portion works by injecting an echo request packet in a LSP and expecting the remote PE endpoint to receive and reply to the packet. The traceroute function works the same as it does for IP where it sends multiple packets with an increasing TTL to let the packet get progressively farther in the LSP path before sending message indication that the TTL has expired.

  [See Pinging LSPs.]

Multicast

- **Support for multicast features (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports the following multicast protocol (IGMP and PIM) features for forwarding IPv4 and IPv6 traffic:
  - Anycast rendezvous point
  - Auto rendezvous point
  - Bidirectional Forwarding Detection (BFD) for PIM
  - IGMP version 1, version 2, and version 3
  - IGMP filter
  - IGMP proxy (relay)
  - IGMP querier
  - IGMP version 1, version 2, and version 3 snooping
  - Multicast Source Discovery Protocol (MSDP)
  - PIM static rendezvous point
  - PIM source-specific multicast (SSM)
  - PIM sparse mode

[See Multicast Overview.]

Routing Protocols

- **Support for Two-Way Active Measurement Protocol (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports Two-Way Active Measurement Protocol (TWAMP). The TWAMP defines a standard for measuring IP performance between two devices in a network. ACX5448 router supports only the reflector side of TWAMP.

  [See Understanding Two-Way Active Measurement Protocol on Routers.]

Routing Policy and Firewall Filters

- **Support for firewall filters and policers (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports configuring firewall filters on packets (families such as bridge domain, IPv4, IPv6, CCC, MPLS) based on packet match conditions. 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.

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

Security

- **Support for secure boot and BIOS (ACX5448)**—Starting with Junos OS Release 18.2R1, a significant system security enhancement, secureboot, has been introduced in ACX5448 router. The secureboot implementation is based on the UEFI 2.4 standard. BIOS in ACX5448 router has been hardened and is responsible for initializing all the components of the router hardware. The following are some of the key functionalities supported by the BIOS in ACX5448 router:

  - Initialization of hardware components
  - Watchdog support
  - Booting the operating system
  - Diagnostics support
  - Secure boot support

  [See Feature Explorer and enter Secure Boot.]

Software Installation and Upgrade

- **Firmware upgrade (ACX6360 Router)**—Starting in Junos OS Release 18.2R1, you can install or upgrade the system firmware on ACX6360 router.

  Install the firmware package by using:

  - **request system firmware add *path/package-name***

  Upgrade an existing firmware, by using any of the following command:

- **request system firmware upgrade pic**
- **request system firmware upgrade cb**
- **request system firmware upgrade re**
- **request system firmware upgrade fpc**

On the ACX6360 line card, you upgrade the following firmware components:

- Uboot—Responsible for loading the operating system on the line card
- FPGA—Controls all functions of the line card

You can also upgrade the following firmware components:

- RE- FPGA—The RE-FPGA is located on the control board and manages board initialization, reboot, and other functions.
- TIC-FPGA—The TIC-FPGA is located on the 8x CFP2 optical port card and manages access to the optical functions.
- FTC FPGA—The FTC FPGA is located on the fan controllers and controls the fan controllers.

- FPD FPGA—The FPD FPGA is located on the LED board and is responsible for the LED board.
- SIB FPGA—The SIB FPGA is located on the SIB and handles the SIBs

Timing and Synchronization

- **Support for PTP transparent clock (ACX5448)**—Starting with Junos OS Release 18.2R1, ACX5448 router supports the PTP transparent clock functionality for PTP over IP, as well as PTP over Ethernet. A certain amount of delay is always experienced by the PTP packets due to queuing and buffering within the router, which could be due to network load or based on the architecture of the router. The PTP transparent clock measures the residence time (the time that the packet spends passing through the router), and adds the residence time into the correction field of the PTP packet. ACX5448 routers support end-to-end transparent clocks. With an end-to-end transparent clock, only the residence time is included in the correction field of the PTP packets. ACX5448 supports end-to-end (e2e) transparent clocks as defined in IEEE1588.

  [See Understanding Transparent Clocks in Precision Time Protocol.]

Port Security

- **MACsec support on ACX6360 routers**—Starting in Junos OS Release 18.2R1, Junos OS supports Media Access Control security (MACsec) on ACX6360 routers. MACsec is an IEEE 802.1AE industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links. MACsec is capable of identifying and preventing most security threats, and can be used in combination with other security protocols such as IP security (IPsec) and Secure Sockets Layer (SSL) to provide end-to-end network security.

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

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 18.2R3 for the ACX Series routers.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (ACX Series)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

Junos OS XML API and Scripting

- **Junos XML protocol <open-configuration> operation no longer emits an uncommitted changes warning (ACX Series)**—Starting in Junos OS Release 18.2R1, the Junos XML protocol **<open-configuration>** operation does not emit an "**uncommitted changes will be discarded on exit**" warning message when opening a private copy of the candidate configuration. However, Junos OS still discards the uncommitted changes upon closing the private copy.
- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (ACX Series)**—Starting in Junos OS Release 18.2R2, 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 3 Features

- **DMA recovery mechanism (ACX Series)**—Starting in Junos OS Release 18.2R1, a potential recovery mechanism has been introduced that is triggered in case the router enters an **Idle** state on any DMA channels. The recovery mechanism resets the necessary registers to recover from failure conditions and therefore a PFE reboot is not required. The following recovery success message is logged in the PFE syslog message:

```
BCM DMA error recovery: Recovery complete Success
```

Network Management and Monitoring

- **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 (ACX Series)**—In Junos OS Release 18.2R1, 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.]

- **Junos OS does not support management of YANG packages in configuration mode (ACX Series)**—Starting in Junos OS Release 18.2R2, 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 18.2R2, 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)**—Starting in Junos OS Release 18.2R2, a recovery mechanism has been introduced that is triggered in case the router enters an Idle state on any DMA channels. The recovery mechanism resets the PFE reboot 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
CHASSISD_IFDEV_DETACH_FPC: ifdev_detach_fpc(0)
```

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 18.2R2, 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

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 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

- OCM 100FX SFPs with this part number are not supported in this release. PR1279202
- With the enterprise profile and multiple masters configured, PTP servo is stuck in FREERUN state after the master fails by disabling the logical interface. PR1281798
- For an ACX5448, the theoretical limit of the ARP learning rate is approximately 150 ARP resolutions per second per logical interface. ARP learning rate is very low and not able to scale to 96000. PR1343221
- L2 rewrite on outgoing MPLS packet is not supported. PR1376001
- Junos OS does not perform VLAN-ID check at the egress; VLAN-ID check is only performed at the ingress. PR1403730

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 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 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 get silently dropped and discarded. PR1194551
- On the ACX5000 platforms, the maximum number of logical interfaces (IFLs) has been increased from 1000 to 4000. PR1229492
- START_BY_START_ERR interrupt handler is not available with the previous version of bcm sdk code. This leads to the status checking of this flag continuously by bcmDPC process, leading to high CPU utilization. As a workaround, add a handler for this interrupt. PR1329656

- The aggregated Ethernet load balancing based on layer 4 information is not working if ports are in different cores of hardware. PR1332448
- 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 aLACP packet comes 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/SrcGport for the untagged/prio-tagged VPLS logical interface. Because the "interface-speicific/input-list" firewall matches SVP/SrcGport in the later stage of pipeline, the LACP packets are not hitting the firewall. PR1346380
- Logical interface classifier information should not be shown in the output of **show class-of-service interface <ifd>** on ACX5000 routers. PR1353828
- On the ACX5000 platforms, after upgrading from Junos OS Release 16.2 and later releases, if the ECC errors occur, the FPC/fxpc process might use high CPU memory. PR1360452
- Remote fault signaling is not supported for 1-Gigabit fiber SFP during autonegotiation. The following cosmetic log errors are observed when the **show interfaces extensive** command is issued: **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
- When PXE boot is used, you might have to create an additional directory grub2/ and move the secure-boot/ folder to the grub2/ directory for PXE image installation to go through. PR1369040
- Because of a race condition, on 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. Further, on deactivating or activating the e1-interface, a core file is generated. PR1378747
- Host-bound traffic might be affected and lt interface might go down in ACX Series routers. PR1382166
- 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 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
- On ACX1000, ACX2000, ACX4000, ACX5048, and ACX5096 routers, after a new child logical interface with VLAN and filter is added on an aggregated Ethernet physical interface or the VLAN ID of a child logical interface with filter is changed. The traffic over the aggregated Ethernet physical interface might get filtered with that filter on the child logical interface. For example, ae-0/0/0 is a physical interface and ae-0/0/0.100 is a logical interface. PR1407855

- On ACX5000 platform, high CPU usage by the fxp process might be seen under a rare condition if parity errors are detected in the devices. It has no direct service/traffic impact. However, because CPU utilization is high during this issue, there are some side effects. For example, it might impact time-sensitive features such as BFD. PR1419761
- 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
- Protocols get forwarded when using non-existing SSM map source address in IGMPv3 instead of pruning. PR1435648

Interfaces and Chassis

- When an unnumbered interface is binding to an interface that has more than one IP address and one of the IPs is deleted, the family inet of the unnumbered interface might be deleted. As a result, traffic loss occur for all the services that rely on the family inet of the unnumbered interface. PR1412534

Layer 2 Features

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

MPLS

- Dynamically configured RSVP LSPs for LDP link protection might not come up after disabling or enabling protocol MPLS. PR1432138

Timing and Synchronization

- Telemetry is not supported on ACX500 line of routers.PR1316570

SEE ALSO

New and Changed Features | 14

Changes in Behavior and Syntax | 25

Known Behavior | 28

Resolved Issues | 32

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

Class of Service (CoS)

- Error message **STUCK_BUFF : port_sp not empty for port 35 sp 1 pkts:1** during the Lag event tests, lag bundle is configured with 64 links. PR1346452

General Routing

- SNMP MIB walk/get/set on jnxDomCurrentTable and jnxDomNotifications might fail on ACX Series platforms. PR1076943
- The 1-gigabit copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. PR1286709
- Port XE-0/3/0 does not come up. PR1328207
- bcmDPC task is high even though interuppt START_BY_START flag set to 0. PR1329656
- On an ACX Series router ring topology, after the link between ACX Series routers and MX Series routers flaps, VPLS RI on the PE device (MX Series) has no MAC of the CE device over L2 circuit. PR1360967
- On ACX5000, **fpc0 (acx_rt_ip_uc_lpm_install:LPM route add failed) Reason : Invalid parameter** is seen after configuring **lpm-profile**. PR1365034

- On ACX5448, channelized ET interface of 25-gigabit interface might not come up after **chassis-control restart**. PR1379288
- The L2 circuit might stop forwarding traffic when one core interface flapping occur. PR1381487
- On ACX Series platforms, the **forwarding-option dhcp-relay forward-only** configuration statement stops working and the DHCP packets are dropped. PR1392261
- MTU is not properly applied and the output of **ping mpls l2circuit sweep** is giving lower values than expected. PR1393947
- ACX5048 rpm rfc2544-benchmarking test fails to start. PR1395730
- Dynamic tunnels are not supported on ACX Series routers. PR1398729
- FPC might crash after offline/online of MIC-3D-16CHE1-T1-CE-H. PR1402563
- VLAN tagged traffic arriving on VPLS interface might get dropped. PR1402626
- On ACX 5448, TrTCM policer configuration parameters are as per RFC4115. PR1405798
- The **show services inline stateful-firewall flow** or **show services inline stateful-firewall flow extensive** command might cause the memory leak. PR1408982
- The ACX Series routers drop DNS responses that contain an underscore. PR1410062
- The aggregated Ethernet interface TWAMP history statistics verification on the client is not as expected, getting **Request Timed Out** error. PR1411344
- VPLS traffic might stop across ACX5000 with the aggregated Ethernet interface. PR1412042
- Junos PCC might reject **PCUpdate** or **PCCreate** message if there is metric type other than type 2. PR1412659
- Number of inet-arp policers implemented on ACX5000 has been increased from 16 to 64. PR1413807
- SWAP memory is not initialized on booting on ACX5048. PR1415898
- Commit error is seen while configuring firewall with the term having log/syslog and accept actions. PR1417377
- CoS table error can sometimes cause traffic outages and SNMP timeouts if the optic is plugged out and inserted back. PR1418696
- In ACX Series platforms, **no-vrf-propagate-ttl** might not work after activating or deactivating of CoS configuration. PR1435791

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: 18.2R2

Class of Service (CoS)

- CoS is incorrectly applied on Packet Forwarding Engine leading to egress traffic drop. PR1329141
- CoS issue with rewrite with certain loss priorities are seen. PR1358721

General Routing

- Incorrect packet statistics are reported in ifHCInUcastPkts OID. PR1306656
- The ACX Series routers supports dual tag to untag traffic Layer3 traffic. PR1307666
- With auto-installation usb configured, the interface related commits might not take effect because of the dcd error. PR1327384
- Memory leak is observed when ACX Series router is under high traffic load. PR1358127
- ACX Series router incorrectly allows to configure higher values in **burst-size-limit** hardware support. PR1361482
- ARP reply is dropped when a temporal buffer-size is added on the NNI interface. PR1363153
- FEC PM error counters are accumulated instead of resetting after bin rollover. PR1363270
- On ACX5000 IPsec SA as OSPFv3 authentication is not working in Junos OS Release 16.2R2 and Junos OS Release 17.3R2. PR1363487
- On ACX5448, **show chassis hardware** shows inconsistent values for PEMs and FANs. PR1364224
- VPLS with **vlan-id-list** not working properly in some releases when PE-CE is aggregated Ethernet interface with single member link and child physical interface flap. PR1365894
- The **commit** or **commit check** might fail because of the error **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- The fxpc might crash after the interface changes on ACX5000. PR1378155
- Per "physical-port-based" filter in the egress firewall. PR1395362

Layer 2 Ethernet Services

- DHCPv6 relay ignores replies from server when renewing. PR1354212

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 for the ACX Series documentation.

SEE ALSO

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 | 14
- Changes in Behavior and Syntax | 25
- Known Behavior | 28
- Known Issues | 29
- Resolved Issues | 32
- Documentation Updates | 35
- Product Compatibility | 37

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://pathfinder.juniper.net/feature-explorer/.

Hardware Compatibility Tool

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

SEE ALSO

New and Changed Features | 14

Changes in Behavior and Syntax | 25

Known Behavior | 28

Known Issues | 29

Resolved Issues | 32

Documentation Updates | 35

Migration, Upgrade, and Downgrade Instructions

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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 18.2R3: EX2300, EX3400, EX4300, EX4600, and EX9200.

Release 18.2R3 New and Changed Features

There are no new features or enhancements to existing features for EX Series in Junos OS Release 18.2R3.

Release 18.2R2 New and Changed Features

Port Security

- **Media Access Control security (MACsec) (EX4300-48MP)—**Starting in Junos OS Release 18.2R2, MACsec is supported on EX4300-48MP switches. MACsec is an 802.1AE IEEE industry-standard security technology that provides secure communication for all traffic on point-to-point Ethernet links.

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

Release 18.2R1 New and Changed Features

Hardware

- **EX4300-48MP and EX4300-48MP-S switches**—Starting with Junos OS Release 18.2R1, two new models of EX4300 switches are available—EX4300-48MP and EX4300-48MP-S switches. These models provide 24 built-in 10/100/1000BASE-T Ethernet network ports, 24 built-in 100/1000/2500/5000/10000BASE-T Ethernet network ports, and four built-in 40-Gigabit Ethernet quad small form-factor pluggable plus (QSFP+) ports that can house 40-Gigabit QSFP+ transceivers. The 24 built-in 10/100/1000BASE-T Ethernet network ports support 10 Mbps, 100 Mbps, and 1 Gbps speeds. The 24 built-in 100/1000/2500/5000/10000BASE-T Ethernet network ports support 100 Mbps, 1 Gbps, 2.5 Gbps, 5 Gbps, and 10 Gbps speeds. All network ports are equipped for PoE+ and provide up to 95 watts of power. The QSFP+ ports are configured as Virtual Chassis Ports (VCPs) by default. You can use them to connect the switches to other devices in a Virtual Chassis configuration.

  [See EX4300 Switch Hardware Guide.]

- **EX9253 switches**—Starting with Junos OS Release 18.2R1, EX9253 switches are available as a modular switch. The switch has two dedicated slots for line cards and supports EX9253-6Q12C and EX9253-6Q12C-M line cards. The switch is available in two variants—with AC power supply and with DC power supply.

  [See EX9253 Switch Hardware Guide.]

Authentication and Access Control

- **Enhancement to NTP authentication method (EX4300)**— Starting in Junos OS Release 18.2R1, Junos OS supports NTP authentication for both SHA-1 and SHA2-256, in addition to the existing NTP authentication method, MD5. You can now choose from among MD5, SHA-1, and SHA2-256 for synchronizing the clocks of Juniper Network routers, switches, and other security devices on the Internet. Using SHA-1 instead of MD5 improves the security of devices with very little impact to timing, while using SHA2-256 provides an increase in security over SHA-1.

  > **NOTE:** By default, network time synchronization is unauthenticated.

  To implement authentication, use **set authentication-key <*key_number*> type** at the **[edit system ntp]** hierarchy level.

  - To enable SHA-1 authentication, use **set authentication key <*key_number*> type sha1 value <*password*>** at the **[edit system ntp]** hierarchy level.
  - To enable SHA2-256 authentication, use **set authentication key <*key_number*> type sha256 value <*password*>** at the **[edit system ntp]** hierarchy level.

[See authentication-key and Configuring NTP Authentication Keys].

Authentication, Authorization, and Accounting (AAA)

- **RADIUS over IPv6 (EX Series)**—Starting with Junos OS Release 18.2R1, EX2300, EX3400, EX4600 and EX4300-48MP switches support IPv6 for user authentication, authorization, and accounting (AAA) using RADIUS servers, in addition to the existing IPv4 support. You can specify which source address Junos OS uses to contact an external RADIUS server. To configure an IPv6 source address for RADIUS authentication, include the **source-address** statement at the [**edit system radius-server *server-address***] hierarchy level. To configure an IPv6 source address for RADIUS accounting, include the **source-address** statement at the [**edit system accounting destination radius server *server-address***] hierarchy level.

  [See source-address.]

Class of Service (CoS)

- **Support for setting unique IEEE 802.1p code point for host-generated RPM packets (EX2300, EX3400, and EX4300)**—You can already set the DSCP code point and IEEE 802.1p code point for all host-generated packets by setting the **dscp-code-point *code-point-value*** option at the **[class-of-service host-outbound-traffic]** hierarchy level, where the first three bits of the defined DSCP code point value are set as the IEEE 802.1p code point value. Starting with Junos OS Release 18.2R1, you can override this IEEE 802.1p code point value for host-generated RPM packets and set a separate value for these packets by setting the **dscp-code-point *code-point-value*** option at the **[services rpm probe *owner* test *test-name*]** hierarchy level, where again the first three bits of the defined DSCP code point value are set as the IEEE 802.1p code point value.

  [See dscp-code-point (Services).]

Dynamic Host Configuration Protocol (DHCP)

- **DHCP smart relay (EX4600)**—Starting with Junos OS Release 18.2R1, you can configure alternative IP addresses for the gateway interface so that if the server fails to reply to the requests sent from the primary gateway address, the switch can resend the requests using alternative gateway addresses. To use this feature, you must configure an IRB interface or Layer 3 subinterface with multiple IP addresses and configure that interface as a relay agent.

  [See Configuring DHCP and BOOTP Relay.]

EVPN

- **NOTE:** NSR and unified ISSU support for point-to-multipoint LSP for EVPN provider tunnel is documented but not supported in Junos OS Release 18.2R1.

  NSR and unified ISSU support for point-to-multipoint LSP for EVPN provider tunnel (EX9200)—Starting in Junos OS Release 18.2R1, 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.

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

- **IGMP snooping support for EVPN-MPLS (EX9200)**—Starting with Junos OS Release 18.2R1, you can configure IGMP snooping on EX9200 switches in an Ethernet VPN (EVPN) over an MPLS network. Enabling IGMP snooping helps to constrain multicast traffic to interested receivers in a broadcast domain.

  Multicast sources and receivers in the EVPN instance (EVI) can each be single-homed to one provider edge (PE) device or multihomed (in all-active mode only) to multiple PE devices. When IGMP snooping is configured with multihomed receivers, IGMP state information is synchronized among peer PE devices by exchanging BGP EVPN Type 7 (Join Sync Route) and Type 8 (Leave Sync Route) network layer reachability information (NLRI). When PE devices receive multicast traffic from the EVPN core on a multihomed Ethernet segment (ES), only the designated forwarder (DF) PE device forwards the traffic, and the DF forwards the traffic only to interested receivers (selective multicast forwarding) based on IGMP snooping reports and BGP EVPN Type 7 routes. PE devices serving single-homed receivers also use selective multicast forwarding based on IGMP snooping reports to forward the traffic only to interested receivers, conserving network bandwidth.

  All PE devices perform inclusive multicast forwarding using ingress replication to forward multicast traffic into the EVPN core to reach all remote PE devices. Multicast traffic at Layer 3 is routed between bridge domains or VLANs using IRB interfaces.

  This feature is supported with multiple EVIs, multicast sources and receivers on the same or different sites, and IGMP snooping in proxy mode only.

  To enable IGMP snooping on PE devices in an EVPN instance, include the **igmp-snooping proxy** statement at the [**edit routing-instances *routing-instance-name* protocols**] or the [**edit routing-instances *routing-instance-name* bridge-domain *bridge-domain-name* protocols**] hierarchy level.

  For inter-VLAN multicast forwarding, PIM distributed DR (PIM DDR) mode must be enabled on all participating IRBs.

  EVPN and IGMP snooping operational mode commands can be used to view information learned from IGMP snooping messages or EVPN Type 7 and Type 8 messages.

  [See Overview of Multicast Forwarding with IGMP Snooping in an EVPN-MPLS Environment.]

- **Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks (EX Series)**—Starting in Junos OS Release 18.2R1, you can configure OSPF, IS-IS, BGP, and static routing with Bidirectional Forwarding Detection (BFD) on an IRB interface that is used as a routed interface in

EVPN. This allows protocol adjacencies to be established between an IRB on a Layer 3 gateway and a CE device and between an IRB on a Layer 3 gateway and a CE device connected to a Layer 2 leaf device in an EVPN-VXLAN network.

[See Supported Protocols on an IRB Interface in EVPN-VXLAN].

- **NOTE:** This feature is documented but not supported in Junos OS Release 18.2R1

EVPN P2MP bud node support (EX9200)—Starting in Junos OS Release 18.2R1, 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].

- **Layer 2 VXLAN gateway in EVPN-VXLAN overlay network (EX4600 switches)**—By using a Layer 3 IP-based underlay network coupled with an Ethernet VPN-Virtual Extensible LAN (EVPN-VXLAN) overlay network, you can deploy larger networks than those possible with traditional Layer 2 Ethernet-based architectures. With overlay networks, endpoints (bare-metal servers [BMSs] or virtual machines [VMs]) can be placed anywhere in the network and remain connected to the same logical Layer 2 network, enabling the virtual topology to be decoupled from the physical topology.

  The physical underlay network over which EVPN-VXLAN is commonly deployed is a two-layer IP fabric, which includes spine and leaf devices. The spine devices provide connectivity between the leaf devices, and the leaf devices function as Layer 2 VXLAN gateways and provide connectivity to the attached endpoints. Starting with Junos OS Release 18.2R1, you can deploy EX4600 switches as leaf nodes in the EVPN-VXLAN overlay network.

  [See Understanding EVPN with VXLAN Data Encapsulation.]

- **EVPN-VXLAN support of Virtual Chassis (EX4600, and EX4600 Virtual Chassis)**—Ethernet VPN (EVPN) supports multihoming active-active mode, which enables a host to be connected to two leaf devices through a Layer 2 LAG interface. Starting with Junos OS Release 18.2R1, the two leaf devices can be EX4600 standalone switches or EX4600 switches configured as a Virtual Chassis.

  On each leaf device, the LAG interface is configured with the same Ethernet segment identifier (ESI) for the host. The two leaf devices on which the same ESI is configured are peers to each other.

[See EVPN-VXLAN Support of Virtual Chassis and Virtual Chassis Fabric].

- **Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network (EX4600 switches)**—Starting in Junos OS Release 18.2R1, EX4600 switches that function as Layer 2 VXLAN tunnel endpoints (VTEPs) can tunnel single-tagged and double-tagged Q-in-Q packets through an Ethernet VPN-Virtual Extensible LAN (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 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:** EX4600 switches do not support the pop and push actions with a configured range of VLANs.

  The ingress and egress VTEPs support the tunneling of Q-in-Q packets and the Q-in-Q actions in the context of specific traffic patterns.

  [See Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network.]

Interfaces and Chassis

- **Support for hyper mode to increase packet processing rate on line cards with enhanced MPCs (EX9200 switches)**—Starting in Junos OS Release 18.2R1, EX9200 line cards that include enhanced MPCs (such as MPC4E and MPC5E) support the hyper mode feature. Enabling the hyper mode feature increases the rate at which a data packet is processed, which results in the optimization of the lifetime of a data packet. Optimization of the data packet lifetime enables better performance and throughput.

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

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

  - Creating Virtual Chassis.
  - Padding Ethernet frames with VLANs.
  - Sending Internet Control Message Protocol (ICMP) redirect messages.
  - Terminating or tunneling subscriber-based services.

[See Understanding the Hyper Mode Feature on Enhanced MPCs for MX Series Routers and EX9200 Switches.]

- **Multi-rate and non-multi-rate support (EX4300-MP switches)**—Starting in Junos OS Release 18.2R1, you can configure an interface to support multiple speeds on EX4300-MP switches. The interfaces now support 2.5G, 5G, and 10G speeds. In previous releases, interfaces supported only 100M and 1G speeds.

  The naming convention for multi-rate interfaces (including 100M and 1G) is “mge-n/n/n”. The differentiation between multi-rate interfaces and 1G interfaces is based on the speed values. The front panel ports have different color coding to differentiate multi-rate and 1G interfaces.

- **4x10SFP+ Uplink Modules support (EX4300-MP Switches)**—Starting in Junos OS Release 18.2R1, you can configure the operating mode on the module to match the type of transceiver you want to use. EX4300-MP switches contain four ports for 10-gigabit small form-factor pluggable (SFP+) transceivers when configured to operate in 10-gigabit mode.

Layer 2 Features

- **L2PT support for tunneling additional protocols (EX2300 and EX3400 switches)**—Starting with Junos OS Release 18.2R1, you can configure Layer 2 protocol tunneling (L2PT) for the following new protocols on EX2300 and EX3400 switches: E-LMI, IEEE 802.1X, MMRP, and UDLD.

  > **NOTE:** Support for tunneling these additional protocols does not apply to multigigabit models of the EX2300 switch (EX2300-24MP or EX2300-48MP).

  [See Layer 2 Protocol Tunneling.]

- **Ethernet ring protection switching (ERPS)(EX2300 and EX3400 switches and Virtual Chassis)**—Starting in Junos OS Release 18.2R1, you can use ERPS to reliably achieve carrier-class network requirements for Ethernet topologies forming a closed loop. ITU-T Recommendation G.8032 version 1 is supported.

  ERPS version 1 comprises the following features:

  - Support for revertive mode of operation of the Ethernet ring
  - Support for multiple ring instances on the same interfaces
  - Support for multiple ring instances on different interfaces
  - Support for interworking with Spanning Tree Protocol, Multiple Spanning Tree Protocol, and redundant trunk groups

  [See Understanding Ethernet Ring Protection Switching Functionality.]

Operation, Administration, and Maintenance (OAM)

- **Ethernet Connectivity Fault Management (CFM) Support (EX2300 and EX3400 switches)**—Starting with Junos OS Release 18.2R1, Connectivity Fault Management (CFM) is supported on EX2300 and EX3400 switches. The major features of CFM are:

- Fault monitoring using the continuity check protocol. This is a neighbor discovery and health check protocol that discovers and maintains adjacencies at the VLAN or link level.
- Path discovery and fault verification using the linktrace protocol. Similar to IP traceroute, this protocol maps the path taken to a destination MAC address through one or more bridged networks between the source and destination
- Fault isolation using the loopback protocol. Similar to IP ping, this protocol works with the continuity check protocol during troubleshooting.

You can configure the Ethernet CFM using the **set protocols oam ethernet connectivity-fault-management** command, and verify the configuration using the **show oam ethernet connectivity-fault-management** command.

- **Ethernet link fault management (LFM) support (EX4600 switches)**—Starting with Junos OS Release 18.2R1, link fault management (LFM) is supported on EX4600 switches. Ethernet OAM provides the tools that network management software and network managers can use to determine how a network of Ethernet links is functioning. The following OAM LFM features are supported:
  - Discovery and link monitoring
  - Remote fault detection

Port Security

- **Media Access Control security with 256-bit cipher suite (EX9200)—**Starting in Junos OS Release 18.2R1, the GCM-AES-256 cipher suite for MACsec in static CAK mode is supported on EX9200 switches with EX9200-40XS line cards installed. The GCM-AES-256 cipher suite has a maximum key length of 256 bits and is also available with extended packet numbering (GCM-AES-XPN-256).

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

- **IP source guard (EX2300 and EX3400 switches and Virtual Chassis)**—Starting with Junos OS Release 18.2R1, you can configure the IP source guard access port security feature to mitigate the effects of source IP address spoofing and source MAC address spoofing. If IP source guard determines that a host connected to an access interface has sent a packet with an invalid source IP address or source MAC address in the packet header, it discards the packet.

  [See Understanding IP Source Guard for Port Security on EX Series Switches.]

- **Support for 802.1X authentication on private VLANs (PVLANs) (EX2300, EX3400, and EX4300 switches and Virtual Chassis)**—Starting in Junos OS Release 18.2R1, you can enable 802.1X (dot1x) authentication for security purposes on access ports that are in a PVLAN.

  PVLANs provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs. PVLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts.

  Authentication prevents unauthenticated devices and users from gaining access to your LAN. For 802.1X and MAC RADIUS authentication, end devices must be authenticated before they receive an IP address from a Dynamic Host Configuration Protocol (DHCP) server.

On a switch that is configured with both 802.1X authentication and PVLANs, when a new device is attached to the PVLAN network, the device is authenticated and then is assigned to a secondary VLAN based on the PVLAN configuration or RADIUS profile. The device then obtains an IP address and is given access to the PVLAN network.

[See Using 802.1X Authentication and Private VLANs Together on the Same Interface.]

- **Private VLANs (EX2300 switches)**—Starting in Junos OS Release 18.2R1, you can enable private VLANs (PVLANs) on EX2300 platforms.

  PVLANs provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs. PVLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts.

  [See Understanding Private VLANs.]

- **Support for DHCP snooping and other access port security features on private VLANs (EX4300 switches and Virtual Chassis)**—Starting in Junos OS Release 18.2R1, you can enable Dynamic Host Configuration Protocol (DHCP) snooping for security purposes on access ports that are in a PVLAN. You can also protect those ports with DHCP options, dynamic ARP inspection (DAI), IP source guard, and neighbor discovery inspection.

  PVLANs provide Layer 2 isolation between ports within a VLAN, splitting a broadcast domain into multiple discrete broadcast subdomains by creating secondary VLANs. PVLANs are useful for restricting the flow of broadcast and unknown unicast traffic and for limiting the communication between known hosts.

  Ethernet LANs are vulnerable to attacks such as address spoofing (forging) and Layer 2 denial of service (DoS) on network devices. The following port security features help protect access ports on your device against loss of information and productivity that such attacks can cause:

  - DHCP snooping—Filters and blocks ingress DHCP server messages on untrusted ports. DHCP snooping builds and maintains a database of DHCP lease information, which is called the DHCP snooping database.
  - DHCPv6 snooping—DHCP snooping for IPv6.
  - DHCP option 82—Also known as the DHCP Relay Agent Information option. Helps protect the switch against attacks such as spoofing of IP addresses and MAC addresses and DHCP IP address starvation.
  - DHCPv6 option 37—Remote ID option for DHCPv6. Used to insert information about the network location of the remote host into DHCPv6 packets.
  - DHCPv6 option 18—Circuit ID option for DHCPv6. Used to insert information about the client port into DHCPv6 packets.
  - DHCPv6 option 16—Vendor ID option for DHCPv6. Used to insert information about the vendor of the client hardware into DHCPv6 packets.
  - DAI—Prevents Address Resolution Protocol (ARP) spoofing attacks. ARP requests and replies are compared against entries in the DHCP snooping database, and filtering decisions are made on the basis of the results of those comparisons.

- IP source guard—Mitigates the effects of IP address spoofing attacks on the Ethernet LAN. The source IP address in the packet sent from an untrusted access interface is validated against the DHCP snooping database.
- IPv6 source guard—IP source guard for IPv6.
- IPv6 neighbor discovery inspection—Prevents IPv6 address spoofing attacks. Neighbor discovery requests and replies are compared against entries in the DHCPv6 snooping database, and filtering decisions are made on the basis of the results of those comparisons.

[See Putting Access Port Security on Private VLANs.]

Restoration Procedures Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (EX Series)**—Starting in Junos OS Release 18.2R1, 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 the 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.]

Software Installation and Upgrade

- **Phone-home client (EX2300 and EX3400 switches)**—Starting with Junos OS Release 18.2R1, 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.

Software Licensing

- **Advanced Feature License (AFL) (EX3400 switches)**—Starting with Junos OS Release 18.2R1, the following features are available as part of the AFL:
  - Border Gateway Protocol (BGP) and multiprotocol BGP (MBGP)
  - IPv6 routing protocols: IPv6 BGP and IPv6 for MBGP
  - IS-IS
  - Virtual routing and forwarding (VRF) BGP

[See Understanding Licenses for EX Series.]

System Management

- **New tool to detect high CPU utilization (EX Series)**—Starting in Junos OS Release 18.2R1, a flight recorder tool is introduced to gather historical data on when the CPU utilization on a device was high and what processes caused the high utilization. The tool collects snapshots of data enabling detection of high CPU usage and faster resolution of issues.

  Because some of the high CPU utilization cases are intentional or expected, you can enable and disable the flight recorder tool to avoid false alarms.

  [See request flight-recorder set high-cpu and show flight-recorder status.]

User Interface and Configuration

- **Support for displaying ephemeral configuration data with filtering (EX Series)**—Starting in Junos OS Release 18.2R1, the **show ephemeral-configuration** command enables you to specify the scope of the configuration data to display. To filter the displayed configuration data, append the statement path of the requested hierarchy to the command.

  [See Displaying Ephemeral Configuration Data in the Junos OS CLI.]

Virtual Chassis

- **Virtual Chassis support (EX4300-48MP)**—Starting in Junos OS Release 18.2R1, EX4300-48MP switches can be interconnected into a Virtual Chassis as one logical device managed as a single chassis. An EX4300-MP Virtual Chassis can contain up to 10 members in either of the following combinations:
  - A non-mixed Virtual Chassis if the members are all EX4300-48MP switches.
  - A mixed Virtual Chassis if the members are a combination of EX4300-48MP switches with other EX4300 switches. The mixed-mode setting is required on all switches. The members in the Routing Engine role must be EX4300-48MP switches, and other EX4300 switches can only be configured in the linecard role. The EX4300-48MP cannot form a mixed Virtual Chassis with any other type of switches.

  The 40-Gbps ports on the rear panel of EX4300-48MP switches are dedicated Virtual Chassis ports (VCPs). You must use those ports to interconnect EX4300-48MP Virtual Chassis members into a non-mixed or mixed Virtual Chassis. The dedicated VCPs cannot be converted into and used as network ports, and no other ports on the EX4300-48MP switch can be used as VCPs. In addition, EX4300 members in a mixed Virtual Chassis with EX4300-48MP members must have a special port mode enabled on VCPs to interconnect with VCPs on EX4300-48MP members. To enable this mode for all VCPs on an EX4300 switch, include the **ieee-clause-82** option when setting mixed mode on the switch, as follows:

  ```
  user@switch> request virtual-chassis mode ieee-clause-82 mixed
  ```

  Otherwise, configuring and administering a non-mixed or mixed mode EX4300-48MP Virtual Chassis is the same as for other EX4300 Virtual Chassis or QFX Series Virtual Chassis.

[See Understanding EX4300 Virtual Chassis.]

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 18.2R3 for the EX Series.

EVPN

- On EX9200 switches, you can configure EVPN to extend a Junos Fusion Enterprise or multichassis link aggregation group (MC-LAG) network over an MPLS network to a data center or campus network. For both Junos Fusion Enterprise and MC-LAG use cases, you must include the **bgp-peer** configuration statement in the **[edit routing-instances *name* protocols evpn mclag]** hierarchy level. This configuration enables the interworking of EVPN-MPLS with Junos Fusion Enterprise or MC-LAG. If you do not include the **bgp-peer** configuration statement in your configuration, unexpected behavior and a core dump could result. To enforce this configuration, we now check for this configuration during the commit. If the configuration is not present, an error occurs.

  See [Understanding EVPN-MPLS Interworking with Junos Fusion Enterprise and MC-LAG .]

General Routing

- **Enhancement to the show interfaces mc-ae extensive command**—You can now view additional LACP information about the LACP partner system ID when you run the **show interfaces mc-ae extensive** command. The output now displays the following two additional fields:
  - Local Partner System ID-LACP partner system ID as seen by the local node.
  - Peer Partner System ID-LACP partner system ID as seen by the MC-AE peer node.

  Previously, the **show interfaces mc-ae extensive** command did not display these additional fields.

  [See show interfaces mc-ae..]

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (EX Series)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

Interfaces and Chassis

- **EEE not supported on mge interfaces operating at 100-Mbps speed (EX2300-24MP and EX2300-48MP)**—In Junos OS Releases 18.1R2, 18.2R1, and later, if both Energy Efficient Ethernet (EEE) and 100-Mbps speed are configured on a rate-selectable (or multirate) Gigabit Ethernet (mge) port on EX2300-24MP and EX2300-48MP switches, the port operates only at 100-Mbps speed but EEE is not enabled on that port. EEE is supported only on mge interfaces that operate at 1-Gbps and 2.5-Gbps speeds.

- **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

- **Junos XML protocol <open-configuration> operation no longer emits an uncommitted changes warning (EX Series)**—Starting in Junos OS Release 18.2R1, the Junos XML protocol **<open-configuration>** operation does not emit an "**uncommitted changes will be discarded on exit**" warning message when opening a private copy of the candidate configuration. However, Junos OS still discards the uncommitted changes upon closing the private copy.
- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (EX Series)**—Starting in Junos OS Release 18.2R2, 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.

Junos Telemetry Interface

- **Change to the configuration location for gRPC-based sensor subscriptions from an external collector (EX Series)**—Starting in Junos OS Release 18.2R1, when an external streaming server, or collector, provisions sensors to export data through gRPC on devices running Junos OS, the sensor configuration is committed to the **junos-analytics** instance of the ephemeral configuration database, and the configuration can be viewed by using the **show ephemeral-configuration instance junos-analytics** operational command. In earlier releases, the sensor configuration is committed to the default instance of the ephemeral configuration database.

Layer 2 Features

- **Configuration option for LLDP VLAN name type, length, and value (TLV) (EX3400, EX4300)**—Starting in Junos OS Release 18.2R1, 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.
- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—From Junos OS Release 18.2R3, the configuration statement **input-native-vlan-push** at the **[edit interfaces**

interface-name] hierarchy level is introduced. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement **input-vlan-map** with a **push** operation is configured.

[See input-native-vlan-push.]

Network Management and Monitoring

- **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 (EX Series)**—In Junos OS Release 18.2R1, 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.]

- **Junos OS does not support management of YANG packages in configuration mode (EX Series)**—Starting in Junos OS Release 18.2R2, 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 18.2R2, 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 Junos OS 18.2R2, 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 Junos OS 18.2R3, 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.

Software Installation and Upgrade

- **New DHCP option introduced for ZTP retry (EX Series)**—Starting in Junos OS Release 18.2R1, a new DHCP option is introduced to set the timeout value for the file downloads over FTP. If the **transfer-mode** is set as FTP, the default value for the time out is automatically set as 120 minutes. That is, if the FTP

session gets interrupted due to loss of connectivity in the middle of a file transfer, it will timeout after 120 minutes and ZTP will attempt to retry the file-fetching process. This value can be overridden using the DHCP option as follows:

```
option NEW_OP.ftp-timeout code 7 = text;
option NEW_OP.ftp-timeout “val”;
```

where **“val”** is the user configurable timeout value in seconds and must be provided (for example, “val”).

Subscriber Management and Services

- **DHCPv6 lease renewal for separate IA renew requests (EX Series)**—Starting in Junos OS Release 18.2R2, 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.]

User Interface and Configuration

- **Changes to the show ephemeral-configuration command (EX Series)**—Starting in Junos OS Release 18.2R1, the **show ephemeral-configuration** operational mode command has the following changes:

  - To display the configuration data in the default instance of the ephemeral configuration database, issue the **show ephemeral-configuration instance default** command. In earlier releases, ephemeral configuration data for the default instance is displayed using the **show ephemeral-configuration** command.
  - To display the configuration data in a user-defined instance of the ephemeral configuration database, issue the **show ephemeral-configuration instance *instance-name*** command. In earlier releases, ephemeral configuration data for a user-defined instance is displayed using the **show ephemeral-configuration *instance-name*** command.
  - To view the complete post-inheritance configuration merged with the configuration data in all instances of the ephemeral database, issue the **show ephemeral-configuration merge** command. In earlier releases, the merged view is displayed using the **show ephemeral-configuration | display merge** command.

- **Change to the maximum number of user-defined instances supported by the ephemeral configuration database (EX Series)**—Starting in Junos OS Release 18.2R1, devices running Junos OS that support configuring the ephemeral configuration database enable configuring a maximum of seven user-defined instances of the ephemeral database. In earlier releases, you can configure up to eight user-defined instances. User-defined instances are configured using the **instance *instance-name*** statement at the **[edit system configuration-database ephemeral]** hierarchy level.

Virtual Chassis

- **New configuration option to disable automatic Virtual Chassis port conversion (EX4300 and EX4600 Virtual Chassis)**—Starting in Junos OS Release 18.2R2, 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 | 38
- Known Behavior | 56
- Known Issues | 59
- Resolved Issues | 66
- Documentation Updates | 79
- Migration, Upgrade, and Downgrade Instructions | 80
- Product Compatibility | 81

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 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.

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

Infrastructure

- OAM Boot Menu options are not supported in EX4300-48MP. PR1336127
- This issue is specific to a downgrade (17.4T) and a core file is seen only once during the downgrade because of a timing issue in the sdk toolkit upgrade. After the upgrade, dcpfe recovers on its own and no issues are seen after that. PR1337008
- When the Layer 3 interface comes up, there can be a mismatch in logical interface counters between the Routing Engine and the Junos telemetry interface. This mismatch pertains to ARP/GARP packets. As the ARP/GARP packets get initiated the moment the Layer 3 interface comes up (from spirent/DUT), Routing Engine ends up with one packet less on the logical interface. PR1361282

Interfaces and Chassis

- 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 was assigned with the identical address after commit. There was no warning during the commit, only syslog messages indicating incorrect configuration. PR1221993
- **EEE not supported on mge interfaces operating at 100-Mbps speed (EX4300-48MP)**—Starting in Junos OS Releases 18.2R1, if both Energy Efficient Ethernet (EEE) and 100-Mbps speed are configured on a rate-selectable (or multirate) Gigabit Ethernet (mge) port, the port operates only at 100-Mbps speed but EEE is not enabled on that port. Note that EEE is supported only on mge interfaces that operate at 1-Gbps, 2.5-Gbps, 5-Gbps, and 10-Gbps speeds.

Platform and Infrastructure

- On EX4300 10-Gigabit links, preexisting MACsec sessions might not come up after the following events: process (pfex, dot1x) restart, system restart, or link flaps. PR1294526
- LAGs with member links of different interface types (for example, ge and mge) is not supported. PR1297309
- 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 the SSD. PR1343091

- On EX4300-48MP, when the primary is corrupted and the switch is power- cycled, the switch gets stuck at Linux after boot. The switch needs to be manually rebooted from the secondary SSD partition and recover the corrupted primary partition. PR1344938
- Broadcast route is not pingable when NTP is configured in broadcast mode. Ping to Broadcast route is not supported. PR1347480
- In case of an aggressive BFD timer value (for example, 1 second), BFD packets get delayed during Virtual Chassis switchover and results in BFD session flap. The minimum BFD timer value should be 3 seconds before Virtual Chassis switchover. PR1356693
- On EX2300 and EX3400 switches, image upgrade might fail due to **insufficient space** issue. PR1376488
- If there are non-recovery (cheap) snapshots present in the system, upgrade may fail due to space constraints and it may be necessary to delete non-recovery (cheap) snapshots to get the upgrade going successfully. PR1470823

Virtual Chassis

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

SEE ALSO

New and Changed Features | **38**

Changes in Behavior and Syntax | **50**

Known Issues | **59**

Resolved Issues | **66**

Documentation Updates | **79**

Migration, Upgrade, and Downgrade Instructions | **80**

Product Compatibility | **81**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 for the EX Series.

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

Authentication and Access Control

- On EX Series switches except EX4300, EX4600, and EX9200, the Link Layer Discovery Protocol (LLDP) core files might be seen when the LLDP neighbor gets expired, and all the information gathered through LLDP will be affected. For example, MAC address and physical layer information, and Power information. PR1408707

General Routing

- There was a timing issue between the Junos OS software and the I2C controllers on an MPC5E during a reboot. The software has been corrected to wait for I2C controllers to be ready before the software starts monitoring the voltage levels and current levels. PR1051902
- Certain QFX and EX Series devices do not pad Ethernet packets with zeros, and thus some packets can contain fragments of system memory or data from the previous packets. This issue is also known as 'Etherleak' and often detected as CVE-2003-0001. Refer to JSA10773 for more information. PR1063645

- 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
- On an EX9200-40XS line card, if you toggle the MACsec encryption option multiple times, encryption and protected MACsec statistics might be updated incorrectly. As a workaround, restart the line card. PR1185659
- The dcpfe process might crash and generate a core file if an unsupported SFP-T is put in the switch. PR1290318
- A loopback filter configured in the switch affects the control traffic from both management and data ports. As a result, the switch management connection is lost, depending on the loopback filter configuration. PR1297264
- In a streaming telemetry scenario, when performing **commit full**, the na-grpd daemon might restart, causing the streaming telemetry to be disconnected. PR1326366
- Default route check for loose mode is added for EX4300-48MP. If a route is taking the default route in loose mode, the packets will be dropped. If that is not set, then any route which is not in the routing table will hit the default route and gets forwarded. PR1341673
- When the primary SSD partition /boot is corrupted and a power-cycle is issued to the switch, the switch will take the default 10 minutes to time out the watchdog timer and then will start booting from the secondary SSD partition. PR1342180
- On the EX4300-48-MP switch, commit of **irb bind** to a VLAN without the IRB logical interface defined is blocked. The IRB logical interface has to be defined before binding the same to the VLAN. PR1342443
- Interface range for channelized interfaces is not supported on EX9253. The user has to configure interfaces individually. PR1350635
- On EX3400, when me0 ports are connected between two EX3400 switches, the link does not come up. The link comes up when me0 is connected to a network port. PR1351757
- The working uplink module SFP-T might go down with Junos OS Release 17.2R1 and later. PR1360602
- On a MACsec static-CAK configuration, the traffic will be blocked expect for STP, Pause, EAPOL, and any other acceptable protocols configured through **exclude-protocol** option. This leads to disruption of all protocols running in the system besides the above mentioned. PR1366031
- 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 drops might be observed with a swapout of a Virtual Chassis of QFX5100 to the EX9253 for testing some heavy multicast traffic, even when IRB interface comes up. PR1369099
- Multicast router advertisement (RA) packets coming on a VLAN need to be flooded to interfaces of all FPCs belonging to the same VLAN. Packets when traversing through HighGig port (that connects different FPCs) need to hit hardware filter to transmit packets in other FPCs. In issue state, the filter is not applicable for the HighGig ports, so multicast RA packets do not traverse through other FPCs. PR1370329

- The interface might not flap when both flap-on-disconnect and port-bounce are sent. PR1372619
- Error messages similar to the following might be observed on MPC cards: **LOG: Err] PQ3_IIC(WR): bus transfer timeout on byte 1 LOG: Err] PQ3_IIC(WR): transfer not complete on byte 1 LOG: Err] PQ3_IIC(WR): I/O error (i2c_stat=0x21, i2c_ctl[0]=0xb0, bus_addr=0x76) LOG: Err] Failed to disable PCA9548(0x76)->channel(0-7) LOG: Err] zlpmb_set_channel: Failed to select channel 0 for MPC-PCIE1V0-LTC3880** One root cause is that, the time to wait for the i2c transaction is not sufficient to finish the i2c transaction intermediately, so once in a while this i2c transaction error is seen. These errors do not impact any functionality on the card. PR1374450
- On EX2300/EX3400 platforms with SFP-T (copper SFP) pluggable module, the interface link status on the SFP-T module might be down while its peer connected interface link status is up. For example, in the EX3400-48T switch with ge-0/2/0 (SFP-T copper SFP) connected back-to-back with ge-0/0/36 (built-in copper port), the link ge-0/2/0 is down, but the link ge-0/0/36 is up. PR1374522
- 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 due to implementation design and chipset limitation. PR1376454
- For EX4300-48MP switches, active SSD firmware upgrade is supported where power-cycle to switch is not required after SSD firmware upgrade. PR1389543
- When the **show** command takes a long time to display results, the STP might change states as BPDUs are no longer processed and cause lots of outages. PR1390330
- If PTP transparent clock is configured on the QFX5200, and if IGMP snooping is configured for the same VLAN as PTP traffic, the PTP over Ethernet traffic might be dropped. PR1395186
- On an EX9200 device with MCLAG configuration and other features enabled, there is a loss of approximately 20 seconds during restart of the routing daemon. This traffic loss varies with the configuration that is done. PR1409773
- The factory-default configuration for EX4300, EX2300, EX3400 and EX4300 MP platforms now include DHCP client configuration on IRB and VME to facilitate connectivity to the phone-home server (redirect.juniper.net) from phone-home-client running on the device. The factory default configuration will include the following:

  dhcp enabled on vme and irb

  default vlan with vlan-id 1 and l3-interface as irb.0PR1423015
- Whenever native VLAN configuration is done along with flexible VLAN tagging on a Layer 3 subinterface, untagged packets will be dropped on that Layer 3 subinterface. PR1434646
- The issue is limited to DB-related to MAC-MOVE scenario. When **dhcp-security** is configured, if multiple IPv4 and IPv6 clients' MAC-MOVE happens, the jdhcpd might consume 100% CPU and jdhcpd will crash afterwards. PR1425206
- Added support for i40e NVM upgrade in EX9208. PR1436223

Infrastructure

- This issue is specific to a downgrade(17.4T) and a core file is seen only once during the downgrade because of a timing issue in the SDK toolkit upgrade after which dcpfe recovers on its own and no issues are seen after that. PR1337008
- The command **request system zeroize** will result in the device going to a continuous reboot on EX platforms. PR1337826
- 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. PR1359339
- 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 the inter-switch link might be dropped. The configuration **inter-switch-link** statement is used when a private PVLAN spans multiple switches. PR1388186

Interfaces and Chassis

- On GRES switchover, VSTP port cost on aggregated Ethernet interfaces might get changed, leading to a topology change. PR1174213

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
- The message **eswd[1200]: ESWD_MAC_SMAC_BRIDGE_MAC_IDENTICAL: Bridge Address Add: XX:XX:db:2b:26:81 SMAC** is equal to bridge mac hence do not learn is seen in syslog every few minutes on the ERPS owner. The logs occur during ERPS PDU in ERPS setup. This message can be ignored. PR1372422

Multicast

- IGMP query packets might be duplicated between Layer 2 interfaces with IGMP snooping enabled. PR1391753

Platform and Infrastructure

- On EX4300, Media Access Control Security (MACsec) might not work properly on PHY84756 1-Gigabit SFP ports, if AN is on and MACsec is configured on those ports. On the EX4300 copper box, all four uplink ports (PIC 2) are attached to PHY84756. On EX4300 fiber box, the last four ports of base board(PIC 0) and eight 1-Gigabit/10-Gigabit uplink ports(PIC 2) are attached to PHY84756. PR1291724
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps happen. Even with enhanced convergence configured, there is no guarantee that subsecond convergence will be achieved. PR1371493
- ICMPv6 packets are hitting the dynamic ingress filter with higher priority, thus never reaching an MF or static classifier. PR1388324
- Adding the second IRB to an aggregated Ethernet and then removing it would cause the first IRB to stop working. PR1423106
- On EX4300 platform with equal-cost multipath enabled, interface flapping might trigger a sequence of ulst next-hop install/uninstall events, which exceed the system limit, leading to next-hop installation failure on the Packet Forwarding Engine. PR1426760

Routing Protocols

- On EX4300 and EX4600 switches, and QFX Series switches (except for QFX10000), if host-destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log/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 GF_DLI (gateway family data links) address, which PIM is unable to process, resulting in the crash. PR1408443

Subscriber Access Management

- Authd reuses addresses too quickly before jdhcpd completely cleans up the old subscribers which causes flooding of error logs 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

VPNs

- MVPN using PIM dense mode does not prune traffic when a Join request is received from an IGMP or PIM client. PR1425876

Known Issues: 18.2R3-S2

- In server_fail scenario, when tagged traffic is sent for first client, MAC learning happen for both data and voice. But for the second client, the same interface learning happen only for voice. Because VLAN is already added for an interface due to first client authentication process. PR1462479
- On EX3400, after loading a scaled configuration, the backup member might rarely crash with a kernel panic. This crash occurs as some of the operations from master member is not synchronized to the backup member. When Virtual Chassis is recovered, system moves to stable state. PR1470163
- On EX9214, after reboot and MACsec enabled link flap, error **errorlib_set_error_log(): err_id(-1718026239)** is observed. PR1448368
- On EX2300 syslog error **jsr_kkcm_socket_accept: soaccept failed** is seen when you commit the configuration. PR1449894
- On a V44 or Junos Fusion Environment system, intermediate traffic drop is seen between AD and SD when sflow is enabled on ingress interface. This is not seen always. When sFLOW is enabled, the original packet is getting corrupted for those packets which hit the sFLOW filter. Because few packets transmitted from the egress of AD1 is short of FCS (4 bytes) + 2 bytes of data due to which the drop occur. It is seen that the normal data packets are of size 128 bytes (4 bytes FCS + 14 bytes Ethernet header + 20 bytes IP header + 90 bytes data) while the corrupted packet is 122 bytes (14 bytes Ethernet header + 20 bytes IP HEADER + 88 bytes data). PR1450373
- If the dynamic assignment of VoIP VLAN is used, the switch might not send correct VoIP VLAN information in LLDP MED packets after any configuration change and commit. PR1458559
- In EX2300 and EX3400 Virtual Chassis while upgrading image through URL option, **/var/tmp/** location might not get cleared automatically and image upgrade might fail due to space constraint in the device. PR1464483
- In a Virtual Chassis, during reboot of one of the members or during mastership switchover, PFEX core file might be generated. PR1465526

- In EX3400 Virtual Chassis during reboot or upgrade, because of a high CPU load in slow path of fxpc TCP keep alive message is not sent. Hence it is observed that some Virtual Chassis members might take longer to join the Virtual Chassis. PR1467707
- On EX3400 traffic loss is seen when SFP-T is connected because of auto-negotiation failure. PR1469750
- With auto-negotiation enabled, EX3400 advertise only 100m whenever we configure the speed 100m. PR1471931
- Under certain conditions, FXPC core files might be generated when renumber FPC master. Traffic might affect by around 1 to 2 minutes. PR1470185
- If Junos OS panics with a filesystem-related panic, such as 'dup alloc', recovery through the OAM shell might be needed. From the OAM shell, run 'fsck' on the root volume until it is marked clean. Only at this point is it safe to reboot to the normal volume. PR1444941
- USB upgrade/recovery might fail with management daemon not responding and unknown class 'junos-login-defaults' when uboot mode date is too far away. PR1454950
- On EX4300 platforms configured with ERP, after multiple devices reboots and then restarts at the same time, ERP might not revert back to the idle state. This issue might be seen in situations where the ERP node-id is not configured manually and after the restart, the default node-id (switch base MAC address) might get reset to 00:00:00:00:00:00, effectively causing multiple devices to have the same node-id. PR1461434
- Though traffic is sent below the configured rate of 80 percent, policing occur because of the bursty traffic and storm in effect messages that are sent to Routing Engine. Burst size allowed 1500 kbs by default and is not user configureable. PR1463979
- On EX4300-MP, incorrect part-number is displayed for SFP+-10G-CU3M dac under **show chassis hardware**. It does not show the complete part number. PR1471583
- On a EX2300 switch, the output of the command **show chassis routing-engine** might display an incorrect value of **mac reset** for the 'last reboot reason' field. PR1331264
- When a file system is full you can expect the system to behave unexpectedly. This results in generating core files like these are bound to happen. PR1450143
- In EX4300 switches when 1-Gigabit Ethernet SFP is connected to 10-Gigabit Ethernet port, auto-negotiation should be disabled (when enabled causes many issues like ARP, link down). Hence when AN is disabled somehow corrupting the TX_DISABLE field hence Laser Tx remain enabled when disabling and plug-out - plug-in.PR1445626
- For a EX4300 system the CLI **set chassis routing-engine on-disk-failure disk-failure-action (reboot | halt)** statement is not supported. PR1450093

SEE ALSO

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

Authentication and Access Control

- Without configuring anything related to dot1x, the syslog message **dot1xd[2192]: task_connect: task PNACAUTH./var/run/authd_control addr /var/run/authd_control: Connection refused** is generated repeatedly. PR1406965

EVPN

- A few minutes of traffic loss might be observed during recovery from link failure. PR1396597
- The device might proxy the ARP probe packets in an EVPN environment. PR1427109

Infrastructure

- IfSpeed and IfHighSpeed erroneously reported as zero on EX2300. PR1326902
- Packet Forwarding Engine is flooded with messages **// pkt rx on ifd NULL unit 0**. PR1381151

Interfaces and Chassis

- The logical interfaces in EVPN routing instances might flap after committing configurations. PR1425339

Layer 2 Features

- On EX2300/EX3400, LLDP packets are dropped at L2PT NNI port when the configuration is applied for the first time. PR1362173

Layer 3 Features

- The l2ald might crash when issuing **clear ethernet-switching table persistent-learning**.PR1381739

Layer 2 Ethernet Services

- The malfunction of the core isolation feature in EVPN-VXLAN scenarios causes traffic to be silently dropped. PR1417729

Network Management and Monitoring

- Over Temperature trap is sent out even though there is Temperature Hot alarm. PR1412161

Platform and Infrastructure

- Ping does not go through the device after WTR timer expires in ERPS scenario. PR1132770
- The RE-PFE out-of-sync errors might be seen in syslog. PR1232178
- OAM Ethernet **connectivity-fault-management** configured on aggregated Ethernet interfaces is not supported but there is no commit error. PR1367588
- IPv6 router advertisement (RA) messages might increase internal kernel memory usage. PR1369638
- Login lockout might never expire because the timestamp of "Lockout start" and "Lockout end" are the same. PR1373803
- RIPv2 update packets might not send with IGMP snooping enabled. PR1375332
- EX4300: upgrade fails during validation of slax script during the upgrade. PR1376750
- ECMP route installation failure with log messages such as **unilist install failure** might be observed on EX4300 device. PR1376804
- Unicast DHCP request might get misforwarded to backup RTG link. PR1388211
- EX3400 Virtual Chassis - **Error tvp_status_led_set** and **Error:tvp_optics_diag_eeprom_read** logs were generated. PR1389407
- **Input rate pps** do not increase on EX2300-MP uplink ports when the packet is a pure Layer 2 packet such as non-etherII or non-EtherSnap. PR1389908
- Continuous log messages get printed on EX4300 after upgrading to Junos OS Release 17.4 or later. PR1391942
- Interface flaps on an EX3400 Virtual Chassis cause interface generated IGMP query packets 224.0.0.1 to be sent to all the members ports, except the master FPC. PR1393405

- PTP over Ethernet traffic could be dropped if IGMP and PTP TC are configured together. PR1395186
- On EX2300 MAC table is not populated after **interface-mode** change. PR1396422
- EX3400 might not learn 30000 MAC addresses when sending MAC learning traffic. PR1399575
- EX4300 OAM LFM might not work on extended VLAN bridge interface with native VLAN configured. PR1399864
- The FBF routing-instance instance-type "forwarding" is missed for EX Series (EX3400). PR1400163
- The authd might crash when issuing **show network-access requests pending** command during the authd restarting. PR1401249
- The TCP connection between ppmd and ppman might be dropped due to a kernel issue. PR1401507
- **adt7470_set_pwm** message is continuously seen after an upgrade to Junos OS Release 18.1R3.3. PR1401709
- The STP does not work when aggregated interfaces number is "ae1000" or above in QFX5000 and "ae480" or above in other QFX or EX Series switches. PR1403338
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface. PR1403528
- 12th and 13th SFP-T ports are going down with the Junos OS Release 18.4R1.3 image installation. PR1404756
- Traffic drop is seen on EX4300 when 10-Gigabit fiber port is using 1-Gigabit Ethernet SFP optics with Auto-Negotiation enabled. PR1405168
- MAC address movement might not happen in Flexible Ethernet Services mode when family inet/inet6 and VLAN bridge are configured on the same physical interface. PR1408230
- EX3400 PSU status is still taking "check" status even though PSU module has been removed. PR1408675
- On EX2300-24P, the error message **dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family**. PR1410717
- The traffic to the NLB server might not be forwarded if the NLB cluster works on multicast mode. PR1411549
- EX Series/QFX Series: PEM alarm for backup FPC will remain on master FPC though backup FPC was detached from Virtual Chassis. PR1412429
- EX4300-48MP: Chassis Status LED shows yellow instead of amber. PR1413194
- EX4300 Q-in-Q - untagged UNI traffic egress as single-tagged on NNI interface. PR1413700
- Chassisd output power budget is received continually per 5 seconds without any alarm after the upgrade to Junos OS Release 18.1R3. PR1414267
- VXLAN Encapsulation next Hop (VENH) does not get installed during BGP flap or restart routing. PR1415450

- EX3400: **show chassis environment** repeats "OK" and "Failed" at short intervals. PR1417839
- The EX3400 Virtual Chassis status might be unstable during the restarting of the Virtual Chassis or after the Virtual Chassis port flaps. PR1418490
- EX4300: Runt counter never incremented. PR1419724
- EX4300 does not send **Fragmentation needed** message when MTU is exceeded with DF bit set. PR1419893
- The pfex process might crash and generate core files when SFP is reinserted. PR1421257
- Virtual Chassis might become unstable and FXPC generates core files when there are a lot of configured filter entries. PR1422132
- Traffic loss is experienced when one of the logical interfaces on LAG is deactivated or deleted. PR1422920
- Multicast traffic might be silently dropped on ingress port with igmp-snooping enabled. PR1423556
- MACsec connection on EX4600 will not come back up after interface disconnect while traffic is passing. PR1423597
- On MX204 optics "SFP-1GE-FE-E-T" I2C read errors are seen when an SFP-T is inserted into a disabled state port. PR1423858
- Auditd crashed when accounting RADIUS server is not reachable. PR1424030
- The native VLAN ID of packets might fail to be removed when forwarded out. PR1424174
- Interface flapping scenario might lead to ECMP nexthop install failure on EX4300. PR1426760
- fxpc core files are generated on EX2300 Virtual Chassis. PR1427391
- Rebooting Virtual Chassis member causes traffic on RTG link to be down for about 30 seconds. PR1427500
- VIP might not forward the traffic if VRRP is configured on an aggregated Ethernet interface. PR1428124
- EX2300-24P: L2Ald core files are observed after removal and readdition of multiple supplicant mode with PVLAN on interface. PR1428469
- Verification of ND inspection with a dynamically bound client, moved to a different VLAN on the same Port is failing. PR1428769
- EX4300 does not drop FCS frames on XE interfaces. PR1429865
- Incorrect model information while polling via SNMP from Virtual Chassis. PR1431135
- The ERPS failover does not work as expected on EX4300 device. PR1432397
- i40e NVM upgrade support for EX platform. PR1436223

Routing Protocols

- The PPM mode for BFD session in EX4300 is centralized and not distributed by default. PR1361800
- EX4300 might drop incoming IS-IS hello packets when IGMP or MLD snooping is configured. PR1400838

- Host-generated ICMPv6 RA packets might be dropped on the backup member of Virtual Chassis if **igmp-snooping** is configured. PR1413543
- The QFX and EX Series switch might not install all IRB MAC addresses in the initialization. PR1416025
- Sometimes, IGMP snooping might not work. PR1420921

Software Installation and Upgrade

- Configuration loss and traffic loss might be seen if the backup Routing Engine is zeroized and is then switched over to master within a short time. PR1389268

Spanning Tree Protocols

- The l2cpd might crash if the VSTP traceoptions and VSTP VLAN all commands are configured. PR1407469

Resolved Issues: 18.2R3-S1

- In EVPN scenario when local L2 interfaces are down, IRB interfaces might also observed to be down even when IM (inclusive multicast) route is available. This might cause a traffic impact. PR1436207
- In an EVPN scenario, if the 25-Gigabit Ethernet interface of Leaf node is configured with an Ethernet Segment Identifier (ESI), and it actually only has a single-homed to reach its peer, that might cause the packets to the peer to be discarded. PR1438227
- On EX2300 as CE/PE device, transit OSPF traffic over Q-in-Q tunneling might be dropped if a firewall filter is applied to Lo0 interface. PR1355111
- l2ald process might crash and generate a core file on EX Series VC when converted a trunk port to dot1x access port while tagged traffic is flowing. There might be a race-condition, where interface mode is being changed while traffic is running and l2ald has processed interface delete but dot1x has not. PR1362587
- On EX Series platform, if storm control is applied on multiple ports, storm control logging might not take effect. PR1401086
- On EX2300 and EX3400 devices, the software installation fails with an error message indicating that there is not enough space to unpack the software image. PR1417441
- If the 2 consecutively produced switches placed in the same Layer 2 network, then their MAC might have overlapped before this fix. PR1425123
- In a Virtual Chassis scenario, when the interfaces flap or VLAN configuration is changed frequently, the network topology will be changed accordingly, then CPU utilization will be dramatically increased to very high within a short time, which might cause the failure of essential communications between VC master and members. When the communication fail, FPC will automatically restart. As a result, VC is split and traffic is lost. PR1427075
- On EX2300/3400 Virtual Chassis platforms in GRES/NSB scenario, if the RSTP/MSTP is enabled, after the shutdown of the master Routing Engine (by 'request system halt' or power shutdown), the GRES is triggered but the delay in transmission of BPDUs might occur for several seconds. Apart from this, if the **bpdu-timeout-action block** statement is enabled on the RSTP/MSTP peer, the STP re-convergence

might occur instead of RSTP/MSTP re-convergence, which results in traffic loss for about 30 seconds. PR1428935

- When the native VLAN is configured along with the flexible VLAN tagging on a L3 subinterface, untagged packets might be dropped on that L3 subinterface. PR1434646
- In a dual mc-ae scenario, if an LACP active device reboots or all AEs are disabled and then enabled on the device, the LACP partner and its mc-ae peer might have different partner system ID, it causes mc-ae to get stuck in waiting state resulting in a traffic impact in the network. PR1435874
- On EX9200s, when configuring too many VLANs and interfaces under VSTP a commit error might occur **xSTP:Trying to configure too many interfaces for given protocol**. PR1438195
- On EX Series platforms with DHCP snooping configuration, the DHCP snooping table of default VLAN ID 1 might be cleared if another VLAN ID is added to the DHCP snooping configuration. The impact is that all the hosts' traffic in the default VLAN 1 might be blocked, especially if other features that leverage the DHCP snooping table (like Dynamic ARP inspection) are also configured on the device. PR1438351
- The rpd process might generate a core file during router boot up due to file pointer issue as there are two code paths that can close the file. We are attempting to close the file without validating the file pointer. PR1438597
- On EX Series next-generation platforms that support "DHCP snooping with PVLAN" (for example, EX4300, EX2300, and EX3400), when using PVLAN with dot1x and dhcp-security, and IRB interface is not configured for the PVLAN, due to the DHCP packets getting dropped on the promiscuous port. Clients in an isolated VLAN might not get IP addresses after completing authentication. PR1442078
- EX3400 FAN alarm (Fan X not spinning) appears and disappears repeatedly after the fantray (absent) is removed. PR1442134
- If DHCPv6 relay is configured, the device might relay the DHCPv6 request without adding "link-layer-type" value to DHCP Option-79 in the relay packet (normally, the value in DHCP option-79 consists of 2 bytes for link-layer type + 6 bytes for client MAC address). When the DHCP server receives this relay packet, it might misunderstand the option value and cannot provide the IPv6 address correctly to the DHCPv6 client. PR1442867
- **/var/host/motd does not exist** message is flooded every 5 seconds in chassisd logs since EX2300 and EX3400 do not support a backup partition. PR1444903
- EX4600 generates a major alarm once any sensor temperature is hit at 56 degrees celsius. PR1446363
- Provisioning an EX4300 device using phone-home client feature can result in a failed upgrade. PR1447291
- On EX3400 platform, because **on-disk-failure** CLI is not supported, when a disk error occurs, the device might go into hang state. For EX3400 virtual chassis, this issue might cause other devices in the VC to stop working. PR1447853
- Version compare in PHC might fail making PHC to download the same image. PR1453535
- On EX4300, when static /64 IPv6 route is configured and points to the interface where uRPF is configured, IPv6 packets which match the routes might be dropped. PR1427866

- On EX4300 Series platforms, the unicast ARP request received might not be replied if **no-arp-trap** option is configured. This can cause ARP resolutions to fail on remote peer devices. PR1429964
- EX4300 has enabled the soft error recovery feature on the Packet Forwarding Engine, which can automatically detect the Packet Forwarding Engine parity error and recover by itself. PR1430079
- On all platforms which support Zero Touch Provisioning (ZTP), the **/var/db/scripts** directory might get deleted after executing **request system zeroize**, and it might not be recreated automatically. PR1436773
- On EX4300 PoE platforms, the PoE might not work if the PoE firmware upgrade hangs (for example, abnormal interruption to the PoE firmware upgrade, such as power failure during upgrade) during the PoE firmware upgrade. As a result, it is unable to provide power to the PoE device. PR1446915
- On EX4300 platform, if FBF filters are applied on IRB with LAG configuration also existing on the box, the firewall filters cannot be created and function correctly due to TCAM programming issues. PR1447012
- Error message **RPD_DYN_CFG_GET_PROF_NAME_FAILED: Get profile name for session XXX failed: -7**, might be seen in syslog after restarting routing daemon. PR1439514
- In the DDOS-protection scenario, when the aggregate bandwidth value (for example, value A) of protocols (l3mtu-fail/ttl/ip-opt/rsvp/ldp/bgp/unknown-l2mc/rip/ospf/stp/pvstp/lldp) is configured, this bandwidth value might be reset to the default value (for example, value B) after the device reboot or Packet Forwarding Engine restart. PR1440847
- On EX Series platforms, the loopback address exported into other VRF instance might not work. PR1449410

Resolved Issues: 18.2R3-S2

- Under EVPN multihoming mode, if ARP request or Neighbor Solicitation (NS) message encapsulated in dual tagged VLAN arrives at the designated forwarder (DF) which might send it back to the local segment as it was, that might cause a loop and at last, overwhelms the device. PR1459830
- On EX2300/EX3400 Virtual Chassis setup, the interface on failed member FPC retains as up state for 120 seconds. This issue might cause traffic loss of about 120 seconds. PR1422507
- With MLD snooping enabled, IPv6 multicast traffic might be dropped on Virtual Chassis if ingress and egress interfaces are on different VC members. PR1423310
- There is a sequence issue when Virtual Chassis member reboots in an aggregated interface. After the VC member reboots, the Routing Engine kernel inject MAC entry to FPC that reboots. Because of the sequence issue, Routing Engine added MAC entry, originally source MAC entry, to FPC as remote MAC entry. And MAC entry is never aged out because it is a remote entry. PR1440574
- After converging VSTP, if there is a VSTP configuration change and then BPDU might not be flooded because of which port role might be in incorrect state in the adjacent switches. There is no loop created in the network. PR1443489
- If a firewall filter is configured with the action 'then vlan' in a VC scenario on some specific platforms (for example, EX2300, EX3400, and EX4600), some of the traffic which matches that filter might be dropped. PR1446844

- When a unicast ARP request is received by EX3400 switch and it is configured with **set switch-options no-arp-trap option**, the ARP request might not be replied. This has been fixed and unicast ARP request will be replied even with **set switch-options no-arp-trap option** configuration. PR1448071
- On EX3400 platform, IPv6 routes received through BGP routing protocol might show an age time of '00:00:00' when displayed using the CLI command **show route**. PR1449305
- From Junos OS release 14.1X53-D15, 15.1R1, and later due to a software defect, DHCP snooping static binding might not take effect after deleting and readding the entries with commit. As a workaround, use **commit full** after the configuration changes. PR1451688
- On EX3400 with half duplex mode on 10M or 100M speed at medium traffic rates, MAC pause frames will be seen on the port and egress traffic on the port will stop to flow. PR1452209
- The VLAN specific parameters might not be used if configuring VLAN all option and VLAN specific configuration. PR1453505
- EX2300 switches generate SNMP trap for high temperature after upgrading to any of the affected Junos OS software. This is due to a temperature threshold value being set incorrectly in the software, SNMP false trap related to temperature gets generated and results in **over temperature** logs. PR1457456
- Storage space limitation leads to image installation failure during Phone home on EX2300 and EX3400 platforms. PR1460087
- With the statement **system ports console log-out-on-disconnect** configured, if executing some operations on console, the console operations might fail to work properly. PR1433224
- On EX2300/EX2300-C platforms, if Junos OS software 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. PR1442376
- Certain EX Series platforms might generate vmcore by panic and gets reset. This is a rare case since it occurs only when Junos FreeBSD Extension statistic- too_long_complete is incremented. PR1456668
- VRRP-V6 state is flapping with init and idle states after configuring **vlan-tagging**. PR1445370
- In a Junos Fusion Enterprise environment, when traffic originates from a peer device connected to the aggregation device and the ICL is a LAG, there might be a reachability issue if the cascade port is disabled and traffic has to flow through the ICL LAG to reach the satellite device. As a workaround, use single interface as the ICL instead of a LAG. PR1447873
- In DHCP relay scenario, if the device (DHCP relay) receives a request packet with option 50 where the requested IP address matches the IP address of an existing subscriber session, such request packet would be dropped. In such a case the subscriber may need more time to get IP address assigned. The subscriber may remain in this state until it's lease expires if it has previously bound with the address in the option 50. PR1435039
- On EX2300, EX3400, EX4300, and EX4600 and QFX Series switches except for QFX10k, if committing the configuration all together (for example, after the reboot), the fxpc/PFE core files might be generated. In the Virtual Chassis scenario, the VC members might be splitted because the VC ports might not be created in time. PR1467763

- On EX2300/EX3400/EX4300/EX4600 platforms, DMA buffer leaking might hit once the next hop of received traffics is not resolved and eventually cause an FPC/pfex crash if the DMA buffer runs exhaustion. PR1436642
- CM errors on certain MPC line cards are classified as major which should be minor or non-fatal. If these errors are generated, it might get projected as a bad hardware condition and therefore triggers Packet Forwarding Engine disable action. PR1449427
- On QFX5100/EX4600 switches due to Bad Chip ID, an fxpc core can be seen during the device reboot. This is due to a transient error related to a chip where vendor tries to get the chip ID and it results in improper info. PR1432023
- On QFX5K/EX4600 with SP (Service Provider) style VLAN configuration (in this method, each VLAN-ID is locally significant to a physical interface), if interface-mac-limit/mac-table-size is configured (i.e. software MAC learning is enabled) and the scale of MAC addresses on the box is more than 2000, traffic might be dropped after QinQ enabled interface is flapped or a change is made to the vlan-id-list. PR1441402
- On EX Series platforms, when there is MAC change for LDP neighbor and IP remains the same, ARP update is proper but MPLS LDP might still use the stale MAC of the neighbor. If there is any application/service such as MP-BGP using LDP as next hop, all transit traffic pointing to the stale MAC will be dropped. PR1451217
- Problem with access to J-web after upgrading Junos OS Release 18.2R2 to 18.2R3, causing incorrect permissions in the php session dir. PR1454150
- Current MAC address might change when deleting one of the multiple L3 interfaces and it has traffic impact when this issue occurs. PR1449206
- On EX2300 and EX3400 platforms, the recovery snapshot might not be able to be created after a system zeroize. This is due to certain hardware space limitation over time where there is not enough space to save full snapshot. PR1439189

Resolved Issues:18.2R2

Authentication and Access Control

- On EX4300-48MP, need to hide commands to configure DHCPv6 client. PR1373691

EVPN

- Proxy ARP might not work as expected in an EVPN environment. PR1368911

General Routing

- EX4300-32F MACsec session stays down on 1G or 10G links after events when events are performed with running traffic. PR1299484
- EX23 and EX34 bridge-id is assigned to "02:00:00:00:00:10" irrespective of base-mac addresses. PR1315633

- Incorrect value of optical power is displayed. PR1326642
- CoS is incorrectly applied on Packet Forwarding Engine, leading to egress traffic drop. PR1329141
- 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
- VLAN change for reauthentication and CoA scenarios is supported only when no other client is authenticated in the same PVLAN domain. PR1346936
- After an FPC becomes online, the other FPC's CPU usage might go up to 100 percent and have traffic loss for around 30 seconds. PR1346949
- The 40G interfaces might not forward traffic. PR1349675
- On EX4300-48MP, when DAI and IPSG are configured for many VLANs, then DAI statistics for one interface show incorrect values. PR1355963
- On EX2300, EX3400, and EX4300MP platforms in a Virtual Chassis setup, dynamic ARP inspection might fail after Virtual Chassis switchover when VSTP is enabled along with **no-mac-table-binding**. PR1359753
- When EX2300/EX3400 platforms are used as transit switches, the traffic sent out of an IRB interface might use the original MAC address instead of the configured MAC address for the IRB interface. PR1359816
- On EX2300MP platforms, the fan count is incorrect in jnxFruName, jnxFilledDescr, and jnxContainersCount.4. PR1361025
- On EX4300-48MP platforms, dot1x protocol subsystem is taking a long time to respond to management requests with the error **the dot1x-protocol subsystem is not responding to management requests**. PR1361398
- Non-existent fan tray 1 reported by chassisd on EX2300. PR1361696
- EX4300MP 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 and MS-MIC are brought offline or online. PR1362734
- FPM board status is missing in SNMP MIB walk result. PR1364246
- On EX2300 platforms, **show filter hw summary** is showing incomplete output. PR1364930
- The l2cpd process might crash when configuring MVRP with private VLAN and RSTP interface all. PR1365937
- Virtual Chassis split followed by fxpc core file might occur upon scaling VLAN members. PR1369678
- Unicast ARP packet loop might be observed in DAI scenario. PR1370607

- NTP broadcast packets are not forwarded out on L2 ports. PR1371035
- MAC refresh packet might not be sent out from the new primary link after the RTG failover. PR1372999
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. PR1373807
- FPC might crash when flapping the output interface of analyzer or sampling. PR1374861
- Port access list group is not properly reallocating TCAM slices. PR1375022
- On EX4300-48MP, syslog error **Error in bcm_port_sample_rate_set(ifl_cmd) : Reason Invalid port** is seen. PR1376504
- The interface ae480 or above might be in STP discarding state on the EX9200. PR1378272
- MACsec issue on EX3400 Virtual Chassis running on Junos OS Release 15.1X53-D59. PR1378710
- ARP request packets might be sent out with 802.1Q VLAN tag even though the outgoing interface is an access port. PR1379138
- On EX4300-48MP, the IP transit traffic hits the lo0 filter. PR1379328
- All interfaces belonging to certain FPCs might be lost after multiple GRES in Virtual Chassis. PR1379790
- The dot1x does not work with Microsoft NPS server. PR1381017
- On EX4300-48MP, the **session-option** stanza under the access profile hierarchy for EX Series and QFX 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

High Availability (HA) and Resiliency

- The backup Routing Engine might go to db prompt after removing or restoring the configuration. PR1269383

Infrastructure

- Unable to provide management when em0 interface of FPC is connected to another FPC Layer 2 interface of the same Virtual Chassis. PR1299385
- The upgrade might fail if bad blocks occur in the flash memory device or file system. PR1317628
- Need support for archiving dmesg file **/var/run/dmesg.boot***. PR1327021
- Enabling **mac-move-limit** stops ping on flexible-vlan-tagging enabled interface. PR1357742
- Core file is generated upon attempt to commit configuration. PR1376362

Interfaces and Chassis

- MC-LAG peer does not send ARP request to the host. PR1360216

Layer 2 Features

- The dcpfe and fxpc process might crash on Packet Forwarding Engines with low memory while allocating a huge memory. PR1362332
- RTG MAC refresh packets will be sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. PR1389695

Platform and Infrastructure

- The mismatch of vlan-id between a logical interface and VLAN configuration might result in traffic getting silently dropped and discarded. PR1259310
- Packet drop might be seen on the lt-x/2/x or lt-x/3/x logical tunnel interfaces. PR1345727
- The ports using SFP-T transceiver might still be up after the system stops. PR1354857
- Interface flapping is seen on EX4300 switch. PR1361483
- Some interfaces cannot be added under MSTP configuration. PR1363625
- On EX4300 and EX4600 platforms, the l2ald process might crash in dot1x scenario. PR1363964
- The Packet Forwarding Engine might crash when it encounters a frequent MAC move. PR1367141
- Forwarding broken after adding protocol EVPN **extended-vlan-id**. PR1368802
- LLDP TLV with incorrect switch port capabilities might be sent. PR1372966
- On EX4300-48MP, the unsupported 1G optics in 10G uplink module cause interface traffic drop. PR1374390
- Packet drop to the router is observed with indirect next hop when load balancing is configured. PR1376057
- Packet drops on interface if the statement **gigether-options loopback** is configured. PR1380746
- IRB interface does not turn down when the master chassis of the Virtual Chassis is rebooted or halted. PR1381272
- 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
- EX4300 device chooses incorrect bridge-id as RSTP bridge-id. PR1383356
- On EX4300-48MP mixed Virutal Chassis, the Packet Ordering Engine interface maximum power configuration on EX4300 member gives error when configured more than 30. PR1383717
- Layer 3 IP route will be destroyed after Layer 2 next-hop change is seen. PR1389688

Routing Protocols

- On EX4300-48MP, stale VLAN entries are seen after continuous script run involving split, merge, and reboot. PR1363739

Resolved Issues: 18.2R1

Forwarding and Sampling

- DHCP service crashes after EX9251 switch is set to factory default by zeroize. PR1329682

General Routing

- Traffic loss is observed while performing NSSU. PR1311977
- The major alarm **Fan and PSU Airflow direction mismatch** might be seen when removing the management cable. PR1327561
- A new configuration statement operational status detail statement is added in **show poe interface**. PR1330183
- The rpd process generates a core file on new backup Routing Engine at task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler after disabling NSR+GRES. PR1330750
- Cannot install backup Linux first when both SSD partitions are corrupted. PR1342168
- On EX2300-24MP chassis, the FAN count is incorrect in jnxFruName, jnxFilledDescr and jnxContainersCount.4. PR1361025
- On EX4300-48MP, while running regression scripts, syslog error **Error in bcm_port_sample_rate_set(ifl_cmd) : Reason Invalid port** is seen. 1376504
- IP transit traffic hits the lo0 filter. PR1379328
- In EX4300-48MP on rare occasion, when **arp-inspection** and **ip-source-guard** are configured for around 150 VLANs together, then some port might show incorrect large value for DAI statistics. PR1379443
- On rare occasions in EX4300-48MP, when **dynamic-arp-inspection**and **ip-source-guard** are removed and added back for around 150 VLANs in one go, then **arp-inspection** statistics for one of the port shows garbage value. PR1379447

Interfaces and Chassis

- On EX2300 and EX3400, IPv6 neighborship is not created on the IRB interface. PR1198482
- On all Junos OS platforms with MC-LAG and VRRP enabled, ARP request might be generated with IRB IP and IRB MAC instead of VIP and VRRP MAC if MC-LAG and VRRP configuration are done in a single commit. PR1257246
- The interface might not work properly after FPC restarts. PR1329896

Layer 2 Ethernet Services

- EX Series platforms might display a false positive CB alarm **PMBus Device Fail**. PR1298612

Layer 2 Features

- The DCPFE/FXPC process might crash and generate a core file. PR1362332

MPLS

- A unified ISSU is not supported with MPLS configuration. PR1264786

Platform and Infrastructure

- Autonegotiation is not working as expected between EX4300 and SRX5800. PR1311458
- The FPC might crash because of the memory leak caused by the VTEP traffic. PR1356279

SEE ALSO

New and Changed Features | **38**

Changes in Behavior and Syntax | **50**

Known Behavior | **56**

Known Issues | **59**

Documentation Updates | **79**

Migration, Upgrade, and Downgrade Instructions | **80**

Product Compatibility | **81**

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R2 documentation for the EX Series switches.

SEE ALSO

New and Changed Features | **38**

Changes in Behavior and Syntax | **50**

Known Behavior | **56**

Known Issues | **59**

Resolved Issues | **66**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

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.

> **NOTE:** NSSU is not supported on EX2300-VC/EX3400-VC from Junos OS Release 15.1X53 to Junos OS Release 18.1R1 or later releases. For example, NSSU is not supported from Junos OS Release 15.1X53-D58 to Junos OS Release 18.1R1 or Junos OS Release 15.1X53-D57 to Junos OS Release 18.2R2.

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 | 38
- Changes in Behavior and Syntax | 50
- Known Behavior | 56
- Known Issues | 59
- Resolved Issues | 66
- Documentation Updates | 79
- Product Compatibility | 81

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 81

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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 release for Junos Fusion Enterprise.

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

Release 18.2R3 New and Changed Features

There are no new or changed features in Junos OS Release 18.2R3 for Junos Fusion Enterprise.

Release 18.2R2 New and Changed Features

There are no new or changed features in Junos OS Release 18.2R2 for Junos Fusion Enterprise.

Release 18.2R1 New and Changed Features

Junos Fusion Enterprise

- **Aggregation device support on EX9253 (Junos Fusion Enterprise)**—Starting with Junos OS Release 18.2R1, EX9253 switches are supported as aggregation devices in a Junos Fusion Enterprise. The aggregation device acts as the single point of management for all devices in the Junos Fusion Enterprise. Junos Fusion Enterprise supports the 802.1BR standard.

  [See Junos Fusion Enterprise Overview.]

- **Junos Fusion Enterprise support for EX4600 switches (Junos Fusion Enterprise)**—Starting with Junos OS Release 18.2R1, you can configure EX4600 switches as satellite devices in a Junos Fusion Enterprise topology. The satellite device in a Junos fusion topology is managed and configured by the aggregation device. Junos Fusion Enterprise uses EX9200 switches in the aggregation device role.

[See Junos Fusion Enterprise Overview.]

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 in Junos OS Release 18.2R3 for Junos Fusion Enterprise.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (Junos Fusion Enterprise)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

SEE ALSO

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 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 Enterprise

- In a Junos Fusion Enterprise, it can 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 18.2R3 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

- JFE : It could take 6 to 30 seconds for the traffic to converge when on the Aggregation device JFE is powered OFF or powered ON PR1257057
- In Junos Fusion Enterprise scenario, Junos fusion is not able to add new satellite devices when MC-LAG is configured on EX platform. PR1374982
- In Junos Fusion Enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, loop-detect feature does not work for ports 0-23, since the loop detect filter is not properly applied. PR1426757

SEE ALSO

New and Changed Features | 83

Changes in Behavior and Syntax | 84

Known Behavior | 85

Resolved Issues | 87

Documentation Updates | 88

Migration, Upgrade, and Downgrade Instructions | 89

Product Compatibility | 94

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

- PoE over LLDP negotiation is not supported on Junos Fusion Enterprise setup PR1366106
- error: peer_daemon: bad daemon: scpd on EX9251 running 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 if issuing "clear ethernet-switching table persistent-learning" command PR1409403
- Extended ports in JFE do not adjust MTU when VoIP is enabled PR1411179
- The traffic might get blackholed in Junos Fusion Enterprise scenario with dual-AD PR1417139

Resolved Issues: 18.2R2

- A satellite device does not recover PoE after the device is offline for more than 10 minutes and rejoins the aggregation device. PR1356478
- In a Junos Fusion Enterprise, the satellite device reboots after an automatic POE firmware upgrade. PR1359065
- In a Junos Fusion Enterprise, the ppm-lite process might generate a core file on the satellite devices. PR1364265
- In a Junos Fusion Enterprise, the scpd process is not running on the EX9251. PR1369646

Resolved Issues: 18.2R1

- Mirrored packets are dropped if analyzer output extended port is reachable via the ICL link. PR1211123
- In a Junos Fusion Enterprise, an scpd core file might be seen on an aggregation device when DACL on dot1x enabled port is installed on a single-homed satellite device. PR1328247
- DHCP security binding entries are not syncronized after the FPC comes offline or online. PR1332828
- In a Junos Fusion Enterprise, there is an issue with 802.1X re-authentication. PR1345365

SEE ALSO

New and Changed Features | 83

Changes in Behavior and Syntax | 84

Known Behavior | 85

Known Issues | 86

Documentation Updates | 88

Migration, Upgrade, and Downgrade Instructions | 89

Product Compatibility | 94

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 for Junos Fusion Enterprise documentation.

SEE ALSO

New and Changed Features | 83

Changes in Behavior and Syntax | 84

Known Behavior | 85

Known Issues | 86

Resolved Issues | 87

Migration, Upgrade, and Downgrade Instructions | 89

Product Compatibility | 94

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 Administration Library.

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-x86-64-18.2R2.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-x86-64-18.2R2.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 from Junos OS Release 18.2

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 Release 18.2R2, follow the procedure for upgrading, but replace the 18.1 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

- New and Changed Features | **83**
- Changes in Behavior and Syntax | **84**
- Known Behavior | **85**
- Known Issues | **86**
- Resolved Issues | **87**
- Documentation Updates | **88**
- Product Compatibility | **94**

Product Compatibility

IN THIS SECTION

- Hardware and Software Compatibility | **94**
- Hardware Compatibility Tool | **94**

Hardware and Software Compatibility

For a complete list of all hardware and software requirements for a Junos fusion for 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://pathfinder.juniper.net/feature-explorer/

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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 18.2R3 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.2R3.

Release 18.2R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.2R2.

Release 18.2R1 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.2R1.

SEE ALSO

- Changes in Behavior and Syntax | 97
- Known Behavior | 97
- Known Issues | 98
- Resolved Issues | 99
- Documentation Updates | 101
- Migration, Upgrade, and Downgrade Instructions | 101
- Product Compatibility | 110

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 18.2R3 or later for Junos Fusion Provider Edge.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (Junos Fusion Provider Edge)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

SEE ALSO

New and Changed Features | 96

Known Behavior | 97

Known Issues | 98

Resolved Issues | 99

Documentation Updates | 101

Migration, Upgrade, and Downgrade Instructions | 101

Product Compatibility | 110

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 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.

Junos Fusion

- In a Junos Fusion with an EVPN solution when an aggregation device loses EVPN connectivity with the rest of the aggregation devices, then LACP over extended ports on this core isolated aggregation device will be brought down until EVPN connectivity is restored. PR1327784
- The queue statistics are not supported for aggregated Ethernet interface. PR1345484

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 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.

Junos Fusion

- When an interface is marked **Loop Detect PDU Error: Detected** in the output of the **show interface** command, the **clear error loop-detect interface** must be executed on all aggregation devices to bring up the interface. PR1327366
- All the FPCs might restart after committing the changes to the VLAN/encapsulation on the extended port if the **per-interface-per-member-link** ingress is configured for sourced routing statistic by using the **set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress** command. PR1392071

SEE ALSO

New and Changed Features | **96**

Changes in Behavior and Syntax | **97**

Known Behavior | **97**

Resolved Issues | **99**

Documentation Updates | **101**

Migration, Upgrade, and Downgrade Instructions | **101**

Product Compatibility | **110**

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

Junos Fusion

- BUM traffic might get dropped on peer fusion aggregation device when the link between satellite device and local aggregated device goes down. PR1384440
- On QFX5110, auto-negotiation is not disabled in hardware after setting **no-auto-negotiation** option in the configuration statement. PR1411852

Resolved Issues: 18.2R2

Junos Fusion

- The laser receive power of the extended ports is higher than the output power of the peer link. PR1358007
- The ppmd process on AD might crash when using authentication key-chain with BFD. PR1375647
- An spmd core file might be seen after executing **request support information** on the aggregation device. PR1375732
- The shutdown of the cascade port might lead to the invalidation of the MPC line card. PR1360876
- QFX satellite device might restart in Junos Fusion solutions when copper SFP is used. PR1369062

Resolved Issues: 18.2R1

Class of Service (CoS)

- Aggregated Ethernet link-protection feature is not supported. PR1355498

Junos Fusion

- Duplicated packets might be received on the multicast downstream devices and multicast receivers. PR1316499
- In Junos fusion, the **show interfaces diagnostics optics satellite** command does not display any outputs. PR1327876
- High IGMP leave latency occurs with IGMP snooping in EVPN. PR1327980
- In Junos Fusion, an aggregate device might show a plus sign (+) sign on the ICL link for a satellite device. PR1335373

SEE ALSO

New and Changed Features | 96

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 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 18.2R3 is different 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 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 highly recommend that you see 64-bit Junos OS software when implementing Junos Fusion Provider Edge.

```
user@host> request system software add validate reboot
source/jinstall64-18.2R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-18.2R3.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

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

```
user@host> request system software add validate reboot
source/jinstall64-18.2R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot
source/jinstall-18.2R3.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 18.2R1 **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 only be converted to SNOS 3.1 and higher.
> - The switch must be either set 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**.

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 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 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

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 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 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 once 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 18.2R1, 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 Release 18.2

To downgrade from Release 18.2 to another supported release, follow the procedure for upgrading, but replace the 18.2 **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

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 | **96**
- Changes in Behavior and Syntax | **97**
- Known Behavior | **97**
- Known Issues | **98**
- Resolved Issues | **99**
- Documentation Updates | **101**
- Migration, Upgrade, and Downgrade Instructions | **101**

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

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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 18.2R3 New and Changed Features

There are no new features in Junos OS Release 18.2R3.

Release 18.2R2 New and Changed Features

Network Management and Monitoring

- **New major alarms on MX Series routers with MPC1 and MPC2**—Starting in Junos OS Release 18.2R2, 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'. By default, these errors result in the Packet Forwarding Engines (PFEs) being disabled. You can use the **show chassis fpc errors** command to view the error details.

  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).

[See show chassis fpc errors.]

Routing Protocols

- **MPLS transit route installation as primary MPLS fast reroute (FRR) for BGP labeled unicast prefixes (MX Series)**—Starting in Junos OS Release 18.2R2, when a peer autonomous system (AS) boundary router or a link fails, traffic traversing through an inter-AS link can be rerouted provided a loop-free path is available. In networks with node protection enabled, MPLS transit routes are installed as primary backup paths for BGP labeled unicast prefixes learned from external BGP multihop sessions. This feature facilitates quicker route resolution and BGP convergence for BGP labeled unicast prefixes.

  To enable node protection in an inter-AS environment for BGP labeled unicast prefixes, include the **protection** configuration statement at the **[edit protocols bgp family]** hierarchy level in **enhanced-ip network-services** mode.

- **Support for BGP routes with N-Multipath primary and 1-Protection backup gateway (MX Series)**—In Junos OS 18.2R2 and later 18.2Rx maintenace releases, the following enhancements are made to the Junos OS:

  - Support N+1 formation for BGP labelled unicast protection (LU).
  - Support N+1 formation for BGP PIC (IPv4, IPv6, LU).
  - Support for hetero-nexthops (ListNH) in such N+1 formations.
  - Support for KRT to defer fib-update if BGP-multipath is in progress.
  - Removed restriction to use **delay-route-advertisement** statement for IPv4 labeled-unicast.
  - Four new options **import**, **install-address <address>**, **no-install**, and **rib (inet.0 | inet6.0)** are added under the **egress-te** statement.
  - A new configuration statement **allow-protection** is introduced to allow protection for multipath legs. To allow protection for multipath legs, use **set allow-protection statement** at the **[edit protocols bgp multipath]** hierarchy level.
  - A new option **always-wait-for-krt-drain** is introduced under **delay-route-advertisement** statement to make more-specific BGP-routes re-advertisement to wait for KRT-queue to drain. To configure this, use **set always-wait-for-krt-drain** at the **[edit protocols bgp family inet unicast delay-route-advertisements]** hierarchy level.

[See allow-protection (Multipath), delay-route-advertisements and egress-te.]

Release 18.2R1-S4 New and Changed Features

Services Applications

- **Inline JFlow support for EVPN traffic (MX10008)**—Starting in Junos OS Release 18.2R1-S4, inline Jflow supports sampling under the bridge family. Inline Jflow monitors traffic hitting the bridge family and reports the necessary fields in either version 9 or IPFIX format.

  A new family **bridge** is introduced under the **[edit forwarding-options sampling instance]** hierarchy that monitors all traffic hitting the VPLS or bridge family.

  [See Understanding Inline Active Flow Monitoring.]

Release 18.2R1-S2 New and Changed Features

Routing Protocols

- **Support for IPv4 VPN unicast and IPv6 VPN unicast address families in BGP (MX Series)**—Starting with Junos OS Release 18.2R1-S2, the following address families are supported to enable advertisement or reception of multiple paths to a destination to or from the same BGP peer, instead of advertising or receiving only the active path to or from the same BGP peer, under the **[edit protocols bgp group *group-name*]** hierarchy:
  - IPv4 VPN unicast (**family inet-vpn**)
  - IPv6 VPN unicast (**family inet6-vpn**)

Release 18.2R1 New and Changed Features

Hardware

- **Support for JNP10K-LC2101 MPC (MX10008)**—Starting in Junos OS Release 18.2R1, Junos OS supports a new fixed-configuration MPC, JNP10K-LC2101. A fixed-configuration MPC does not contain separate slots for Modular Interface Cards (MICs). MX10008 routers support eight JNP10K-LC2101 MPCs. The JNP10K-LC2101 MPC provides a maximum bandwidth of 2.4Tbps and has six Packet Forwarding Engines, each providing a maximum bandwidth of up to 400 Gbps, which cannot be oversubscribed. You can configure the bandwidth of the MPC to provide a decreased bandwidth of 1.44Tbps as well, if required. Use the **set chassis fpc *fpc-slot-number* pfe-bandwidth 240g** to modify the forwarding capacity of each Packet Forwarding Engine to 240 Gbps.

  JNP10K-LC2101 supports:

- Multi-rate ports. The ports on the JNP10K-LC2101 MPC support multiple port speeds such as 10 Gbps, 40 Gbps, and 100 Gbps. Hence, they are known as multi-rate ports. All ports support all port speeds. To view the port speed information for each port, use the **show chassis pic fpc-slot *fpc-slot-number* pic-slot *pic-slot-number*** command.
- PIC-based tunnel configuration.
- Maximum transmission unit (MTU) size of 16,000 bytes for transit traffic.
- Dynamic Power Management for effective utilization of available power.
- Flexible queuing supports 128,000 queues per line card, including queues on both ingress and egress interfaces. You can use an additional license to support up to 256,000 queues or 1,500,000 queues per slot.

[See JNP10K-LC2101 MPC on MX10008 Routers Overview.]

- **JNP10K-LC2101 MPC (MX10008)**—Starting with Junos OS Release 18.2R1, JNP10K-LC2101 MPC is supported on the MX10008 router. The JNP10K-LC2101 MPC has fixed MPC ports with 2.4 Tbps and supports 24 100-Gigabit Ethernet QSFP28 ports, and 24 40-Gigabit Ethernet QSFP ports, and 96 10-Gigabit Ethernet ports using a breakout cable (4x10 Gigabit Ethernet). The MPC also supports combinations of 100-Gigabit Ethernet, 40-Gigabit Ethernet, and 10-Gigabit Ethernet ports.
- **New Routing and Control Board REMX2008-X8-128G (MX2008)**—Starting in Junos OS Release 18.2R1, the Routing and Control Board (RCB), REMX2008-X8-128G is supported on MX2008 routers. The RCB has increased memory and storage to support node virtualization . The RCB is equipped with an 8-Core 2.3-GHz processor, 128-GB memory, and two 200-GB SSDs and also supports Secure Boot for enhanced boot security.

  [See MX2008 Routing and Control Board (MX2008 RCB) Description.]

Authentication and Access Control

- **Enhancement to NTP authentication method (MX240, MX480, MX960, MX2020, MX2010)**—Starting in Junos OS Release 18.2R1, Junos OS supports NTP authentication for both SHA-1 and SHA2-256, in addition to the existing NTP authentication method, MD5. You can now choose from among MD5, SHA-1, and SHA2-256 for synchronizing the clocks of Juniper Network routers, switches, and other security devices on the Internet. Using SHA-1 instead of MD5 improves the security of devices with very little impact to timing, while using SHA2-256 provides an increase in security over SHA-1.

  > **NOTE:** By default, network time synchronization is unauthenticated.

  To implement authentication, use **set authentication-key *key_number* type** at the **[edit system ntp]** hierarchy level.

  - To enable SHA-1 authentication, use **set authentication key *key_number* type sha1 value *password*** at the **[edit system ntp]** hierarchy level.

- To enable SHA2-256 authentication, use **set authentication key *key_number* type sha256 value *password*** at the **[edit system ntp]** hierarchy level.

[See authentication-key and Configuring NTP Authentication Keys]

Authentication, Authorization, and Accounting (AAA) (RADIUS)

- **Existing TACACS+ behavior is made VRF aware (MX Series)**—Starting in Junos OS Release 18.2R1, the **routing-instance** statement at the **[edit system tacplus-server *server-address*]** hierarchy level and **[edit system accounting destination tacplus server *server-address*]** hierarchy level can now be used to configure any routing instance present at the **[edit routing-instances]** hierarchy level. TACACS+ traffic uses the configured routing instance. In previous releases, the **routing-instance** statement at the **[edit system tacplus-server *server-address*]** hierarchy level and **[edit system accounting destination tacplus server server-address]** hierarchy level can be used only to configure the mgmt_junos routing instance.

  [See Configuring TACACS+ Authentication and Configuring TACACS+ System Accounting.]

Class of Service (CoS)

- **Support for collecting aggregate queue statistics for underlying logical interfaces (MX Series)**—By default, to preserve memory resources, aggregate queue statistics are not collected for underlying logical interfaces (Level 2 interfaces). Queue statistics are collected for the upper-level logical interfaces (Level 3 and above) and the physical interface (Level 1). By default, the command **show interfaces queue *interface-name*** shows all zeros for underlying logical interfaces. Starting with Junos OS Release 18.2R1, you can enable the collection of aggregate queue statistics for all underlying logical interfaces on a particular physical or aggregate (for example, ae0) interface by including the **logical-interface-aggregate-statistics** statement at the **[edit class-of-service interfaces *interface-name*]** hierarchy level. You can show the aggregate queue statistics by running the **show interfaces queue *interface-name*** command.

  [See logical-interface-aggregate-statistics.]

- **Support for excluding the overhead bytes from queue statistics (MX Series)**—By default, the Layer 2 header bytes applied to upper-level logical interfaces are included in CoS per-queue statistics at the physical interface, which can provide inaccurate results. Starting with Junos OS Release 18.2R1, you can exclude the counting of overhead bytes from aggregate queue statistics by enabling the **exclude-queue-overhead-bytes** option at the **[edit class-of-service interfaces *interface-name*]** hierarchy level. To also exclude the counting of overhead bytes from aggregate queue statistics of all child interfaces, including logical interfaces and interface sets, add the **include-hierarchy** option at the **[edit class-of-service interfaces *interface-name* exclude-queue-overhead-bytes]** hierarchy level.

  [See exclude-queue-overhead-bytes.]

- **Support for bypass-queuing-chip option (vMX)**—Starting with Junos OS Release 18.2R1, the **bypass-queuing-chip** option at the **[edit class-of-service interfaces *interface-name*]** hierarchy level is supported on vMX virtual routers. Enable this option on vMX routers to save a vCPU when scheduling is not needed on an interface. With this option, be careful not to oversubscribe the interface bandwidth.

  [See bypass-queuing-chip.]

EVPN

- > **NOTE:** This feature is documented but not supported in Junos OS Release 18.2R1

  NSR and unified ISSU support for point-to-multipoint LSP for EVPN provider tunnel (MX Series and vMX)—Starting in Junos OS Release 18.2R1, 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 EVPN-VPWS flexible cross-connect (MX Series)**—Starting with Junos OS Release 18.2R1, Ethernet VPN (EVPN) virtual private wire service (VPWS) flexible cross-connect is introduced to address a label resource issue that could occur on some low end access routers. This is possible when there are a group of attachment circuits (ACs) under the same EVPN instance (EVI) and share the same label.

  > **NOTE:** The label resource issue is applicable to a service edge router that is interoperable with the access router that uses FXC scheme to conserve its label usage. It is assumed that the label resource issue does not apply to Juniper Networks service edge router, the vMX (or MX), that uses the pseudowire subscriber interface. Thus there is no change for the label assign scheme on the service edge router with regular EVPN-VPWS pseudowire subscriber head-end termination.

  [See Overview of Flexible Cross-Connect Support on VPWS with EVPN.]

- **Support for head-end termination for EVPN VPWS for business services (MX Series)**—Starting with Junos OS Release 18.2R1, Ethernet VPN (EVPN) virtual private wire service (VPWS) is supported on pseudowire subscriber logical interface.

  Prior to Junos OS 18.2 Release, pseudowire subscriber logical interface is used with either Layer 2 circuit or Layer 2 VPN for psuedowire headend termination service.

  An Ethernet VPN (EVPN) enables you to connect dispersed customer sites using a Layer 2 virtual bridge. Virtual private wire service (VPWS) Layer 2 VPNs employ Layer 2 services over MPLS to build a topology of point-to-point connections that connect end customer sites in a VPN. EVPN-VPWS as a next generation of pseudowire subscriber interface technology brings the benefit of EVPN to point-to-point service by providing fast convergence upon node failure and link failure through its multihoming feature. As a result,

you can use EVPN-VPWS on pseudowire subscriber interface for head-end termination into different services.

You can configure the pseudowire subscriber logical interface for EVPN-VPWS so that the pseudowire established by EVPN-VPWS can be headend terminated into either Layer 3 VPN or BGP-VPLS. The head-end termination covers single (single-homed) pseudowire termination and redundant (multihomed) pseudowire termination into Layer 3 VPN and BGP-VPLS.

[See Overview of Pseudowire Subscriber Logical Interface Support on VPWS with EVPN.]

- **EVPN pure type-5 route support (MX Series)**—Starting with Junos OS Release 18.2R1, you can configure pure type-5 routing in an Ethernet VPN (EVPN) Virtual Extensible LAN (VXLAN) environment. Pure type-5 routing is used when the Layer 2 domain does not exist at the remote data centers. A pure type-5 route advertises the summary IP prefix and includes a BGP extended community called a router MAC, which is used to carry the MAC address of the sending switch and to provide next-hop reachability for the prefix. To configure pure type-5 routing, include the **ip-prefix-routes advertise direct-nexthop** statement at the **[edit routing-instances *routing-instance-name* protocols evpn]** hierarchy level. To enable two-level equal-cost multipath (ECMP) next hops in an EVPN-VXLAN overlay network, you must also include the **overlay-ecmp** statement at the **[edit forwarding-options vxlan-routing]** hierarchy level.

  [See Understanding EVPN Pure Route Type-5.]

- **IGMP snooping support for EVPN-MPLS (MX Series, vMX)**—Starting with Junos OS Release 18.2R1, you can configure IGMP snooping on MX Series routers with MPCs and vMX routers in an Ethernet VPN (EVPN) over an MPLS network. Enabling IGMP snooping helps to constrain multicast traffic to interested receivers in a broadcast domain.

  Multicast sources and receivers in the EVPN instance (EVI) can each be single-homed to one provider edge (PE) device or multihomed (in all-active mode only) to multiple PE devices. When IGMP snooping is configured with multihomed receivers, IGMP state information is synchronized among peer PE devices by exchanging BGP EVPN Type 7 (Join Sync Route) and Type 8 (Leave Sync Route) network layer reachability information (NLRI). When PE devices receive multicast traffic from the EVPN core on a multihomed Ethernet segment (ES), only the designated forwarder (DF) PE device forwards the traffic, and the DF forwards the traffic only to interested receivers (selective multicast forwarding) based on IGMP snooping reports and BGP EVPN Type 7 routes. PE devices serving single-homed receivers also use selective multicast forwarding based on IGMP snooping reports to forward the traffic only to interested receivers, conserving network bandwidth.

  All PE devices perform inclusive multicast forwarding using ingress replication to forward multicast traffic into the EVPN core to reach all remote PE devices. Multicast traffic at Layer 3 is routed between bridge domains or VLANs using IRB interfaces.

  This feature is supported with multiple EVIs, multicast sources and receivers on the same or different sites, and IGMP snooping in proxy mode only.

  To enable IGMP snooping on PE devices in an EVPN instance, include the **igmp-snooping proxy** statement at the [**edit routing-instances *routing-instance-name* protocols**] or the [**edit routing-instances *routing-instance-name* bridge-domain *bridge-domain-name* protocols**] hierarchy level.

For inter-VLAN multicast forwarding, PIM distributed DR (PIM DDR) mode must be enabled on all participating IRBs.

EVPN and IGMP snooping operational mode commands can be used to view information learned from IGMP snooping messages or EVPN Type 7 and Type 8 messages.

[See Overview of Multicast Forwarding with IGMP Snooping in an EVPN-MPLS Environment.]

- **NOTE:** This feature is documented but not supported in Junos OS Release 18.2R1

Support for mLDP P2MP tunnels with EVPN for BUM traffic (MX Series and vMX)—Although present in the code, the ability to configure and signal a P2MP LSP for the EVPN inclusive provider tunnel for BUM traffic is not supported in Junos OS Release 18.2R1. P2MP LSPs manages efficient core bandwidth utilization because it uses multicast replication only at the required nodes instead of ingress replication at the ingress PE node.

- **Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks (MX Series and vMX)**—Starting in Junos OS Release 18.2R1, you can configure OSPF, IS-IS, BGP, and static routing with bidirectional forwarding detection (BFD) on an IRB interface that is used as a routed interface in EVPN. This allows protocol adjacencies to be established between an IRB on a Layer 3 gateway and a CE device connected directly to a Layer 3 gateway or to a Layer 2 leaf device in an EVPN-VXLAN network.

  [See Supported Protocols on an IRB Interface in EVPN-VXLAN .]

- **Support for groupVPN failover to backup router (MX Series and vMX)**—Group VPN is a trusted group to eliminate point-to-point tunnels and their associated overlay routing. All group members share a common security association (SA), known as a group SA (GSA). The GSA enables group members to decrypt traffic that was encrypted by any other group member. Starting in Junos OS Release 18.2R1, Junos OS confirms the Group VPN redundancy with service redundancy daemon running on MX Series routers. MX Series routers with redundancy between them act as Group VPN members.

- **Support for passing of traffic during a policy mismatch between key server and group member (MX Series and vMX)**—Currently, packets that do not match the traffic policy provided by the group key server are dropped by default. Starting in Junos OS Release 18.2R1, you have an option to change the default behavior to disable encryption and forward the packets instead of dropping them.

  You can configure the **forward-policy-mismatch** within the **group vpn object** configuration to enable the support for forwarding policy-mismatched packets at the **[edit security group-vpn member ipsec]** hierarchy level.

- **NOTE:** This feature is documented but not supported in Junos OS Release 18.2R1

EVPN P2MP bud router support (MX Series and vMX)—Starting in Junos OS Release 18.2R1, 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]

Flow-Based and Packet-Based Processing

- **Support for inline flow monitoring (MX10008)**—Starting in Junos OS Release 18.2R1, Junos OS supports inline active flow monitoring. Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Version 9 template is supported for IPv4, IPv6, MPLS, and MPLS-IPv4. IPFIX template is supported for IPv4, IPv6, MPLS, MPLS-IPv4, and VPLS flows. Both IPFIX and version 9 templates use UDP as the transport protocol.

  [See Inline Active Flow Monitoring.]

High Availability (HA) and Resiliency

- **Resiliency support for JNP10K-LC2101 MPC (MX10008)**—Starting in Junos OS Release 18.2R1, resiliency support is enabled for JNP10K-LC2101 MPC on MX10008 routers.

Interfaces and Chassis

- **Enhancement to increase the threshold of corrected single-bit errors (MPC7E, MPC8E, MPC9E on MX Series)**—In Junos OS Release 18.2R1, the threshold of corrected single-bit errors is increased from 32 to 1024, and the alarm severity is changed from Major to Minor for those error messages. There is no operational impact upon corrected single-bit errors. Also, a log message is added to display how many single-bit errors have been corrected between the reported events as follows:

  EA[0:0]: HMCIF Rx: Link0: Corrected single bit errordetected in HMC 0 - Total count 25

  EA[0:0]: HMCIF Rx: Link0: Corrected single bit errordetected in HMC 0 - Total count 26

  [See Alarm Overview.]

- **Fabric management support (MX10008)—**Starting in Junos OS Release 18.2R1, fabric management is supported on MX10008 routers. The fabric architecture of MX10008 routers consists of six Switch Fabric Boards (SFBs). The MX10008 MPC has six Packet Forwarding Engines, each having 24 connections to the fabric (24 fabric planes, or 4 connections per SFB). The MX10008 will have 24 planes active when all the six SFBs are populated. However, in case of a failure of one SFB, line rate can still be achieved with 20 planes (that is, a minimum of five SFBs are required to achieve line rate). The fabric supports a link speed of 25 Gbps. Fabric management involves training the fabric links, monitoring the links, and collecting fabric statistics. The MX10008 also supports fabric hardening.

  [See Fabric Plane Management on JNP10K-LC2101 Overview]

- **Support for FRU control, power management, and environmental monitoring (MX10008)**—Starting with Junos OS Release 18.2R1, Junos OS chassis management software for MX10008 routers with

JNP10K-LC2101 MPC provides enhanced environmental monitoring and FRU control. MX10008 has a pair of Routing Engines, which support virtualization. Each Routing Engine board is a single FRU. All FRUs including Routing Engines, Packet Forwarding Engines, interfaces, power supplies, and fan trays are upgradable. The MX10008 chassis supports two kinds of power supply modules (PSM)—a DC PSM and an AC PSM. The AC PSM delivers 2700 W of power, while the DC PSM delivers 2500 W. The MX10008 cooling system contains two fan trays, with 11 fans in each fan tray. MX10008 supports temperature thresholds for each temperature sensor, which enables the router to precisely control the cooling, raise alarms, and shut down a FRU. The router also supports preserving power-on sequence for the FPCs.

[See Understanding How Dynamic Power Management Enables Better Utilization of Power.]

- **Support for inline Two-Way Active Measurement Protocol (TWAMP) server and client on MX10008—** Starting in Junos OS Release 18.2R1, Mx10008 router with MPC7E cards support the inline Two-Way Active Measurement Protocol (TWAMP) control-client and server for transmission of TWAMP IPv4 UDP probes between the session-sender (control-client) and the session-reflector (server). The TWAMP control-client and server can also work with a third-party server and control-client implementation.

  TWAMP is an open protocol for measuring network performance between any two devices that support TWAMP. To configure the TWAMP server, specify the logical interface on the service PIC that provides the TWAMP service by including the twamp-server statement at the: **[edit interfaces si-fpc/pic/ port unit logical-unit-number rpm]** hierarchy level.

  To configure the TWAMP client, include the twamp-client statement at the: **[edit interfaces si-fpc/pic/ port unit logical-unit-number rpm]** hierarchy level.

  [See *Understanding Two-Way Active Measurement Protocol on Routers*.]

- **Software support for MX10008—**Starting in Junos OS Release 18.2R1, MX10008 routers support the following software features:

  - Class of services (CoS)—Helps prioritize packets to avoid random loss of data when a network experiences congestion and delay.
  - Tunneling and encryption—Encapsulates arbitrary packets inside a transport protocol; and thereby provides a private, secure path through an otherwise public network.
  - Firewall filters—Provide rules that define whether to accept or discard packets that are transiting an interface.
  - Port mirroring—Enables you to analyze traffic on routers and switches that, unlike hubs, do not broadcast packets to every port on the destination device.
  - OpenConfig—Supports the use of vendor-neutral data models to configure and manage the network.
  - Detection of wedge condition—Detects several types of wedge conditions. A wedge condition is caused by an error that blocks network traffic.
  - Junos Telemetry Interface (JTI)—Enables you to provision sensors to collect and export data for various system resources, such as physical interfaces and firewall filters.

- **Support for inline MAP-E BR solution (MX Series routers with MPC and MIC interfaces)**—Starting in Junos OS Release 18.2R1, you can configure Mapping of Address and port – Encapsulation (MAP-E) as an inline service on MX Series routers that use MPC and MIC interfaces. MAP-E is an automatic tunneling mechanism that encapsulates IPv4 packets within an IPv6 address. The IPv4 packets are carried in an IPV4-over-IPV6 tunnel from the MAP-E Customer Edge (CE) devices to the MAP-E Provider Edge (PE) devices (also called as Border Relay (BR) devices) through an IPV6 routing topology, where they are de-tunneled for further processing.

  The MAP-E feature is beneficial for service providers for providing IPv4 connectivity to their subscribers over the ISP's IPv6 access network.

  [See Configuring Mapping of Address and Port with Encapsulation (MAP-E).]

- **Limited encryption Junos OS image and boot restriction (MX2008)—**Starting with Junos OS Release 18.2R1, the MX2008 routers with the Routing Engines REMX2008-X8-64G-LT support only the Junos Limited image. The Junos Limited image does not have data-plane encryption and is intended only for countries in the Eurasian Customs Union because these countries have import restrictions on software containing data-plane encryption. Unlike the Junos Worldwide image, the Junos Limited image supports control-plane encryption through SSH and SSL, thus allowing secure management of the system. The Routing Engines are restricted to boot only the Junos Limited image.

  [See Junos OS Editions.]

- **Support for secure boot and upgraded SSD size and RAM size (MX2008)—**Starting in Junos OS Release 18.2R1, a significant system security enhancement, secure boot, has been introduced. 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 boot loader, and the kernel are cryptographically protected. No action is required to implement secure boot.

  The SSD size and the RAM size of the REMX2008-X8-128G-S Routing Engine is upgraded to 2x200-GB and 128-GB, respectively.

  [See Feature Explorer and enter **Secure Boot**.]

- **Upgraded SSD size and RAM size on the REMX200-8-X8-128G-S Routing Engine (MX2008)—** Starting in Junos OS Release 18.2R1, the SSD size and the RAM size of the REMX2008-X8-128G-S Routing Engine are upgraded to 2x200-GB and 128-GB, respectively. The increased SSD size facilitates increased storage of core and log files.

  [See Salient Features of the Routing Engines with VM Host Support.]

- **Support for 240-V high-voltage DC (HVDC) PSMs and PDMs (MX2008, MX2010, MX2020)**—Starting in Junos OS Release 18.2R1, Junos OS supports 240-V HVDC power supply modules (PSMs; model number: MX2K-PSM-DC-240V) and power distribution modules (PDMs; model number: MX2K-PDM-DC-240V) on the MX2000 line of routers. The PDM supplies 240-V HVDC power to each PSM. The 240-V HVDC power supplies are similar in functionality and physical specifications to the existing DC PSMs and PDMs supported on the MX2000 routers, except that the 240-V HVDC PSMs

and PDMs support 240-V input voltage feed. The 240-V HVDC PSMs and PDMs are supported in HVDC environments that support an input voltage range of 190 VDC through 290 VDC.

- **Support for PTP over Ethernet and hybrid mode over link aggregation group (MX240, MX480, MX960, MX2010, MX2020)**—Starting in Junos OS Release 18.2R1, the MPC5E and MPC6E 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]

Junos Telemetry Interface

- **Streaming OpenConfig data from Routing Engine (RE) sensors over UDP in protobuf format (MX)**–Starting in Junos OS Release 18.2R1, you can stream OpenConfig-based sensor data from Routing Engine sensors using the Junos Telemetry Interface (JTI). This allows you to stream the OpenConfig sensor data in gRPC/protobuf format rather than in key/value pairs. This makes the messages smaller and is more efficient.

  [See Overview of the Junos Telemetry Interface.]

- **Routing Engine state sensors for the Junos Telemetry Interface (MX Series)**—Starting with Junos OS Release 18.2R1, you can export statistics for the Routing Engine state through the Junos Telemetry Interface using the following resource paths:

  - **/junos/kernel-ifstate/stats/churn-rate**
  - **/junos/kernel-ifstate/stats/peer-consumption-rate**
  - **/junos/kernel-ifstate/stats/vetos-statistics**

  Only gRPC streaming is supported.

  To provision the sensor to export data through 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.

  Support for the Junos Telemetry Interface was introduced on QFX10000 and QFX5200 switches in Junos OS Release 17.2R1.

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

- **Expanded ON_CHANGE support for Junos Telemetry Interface (JTI) (MX Series, PTX Series)**—Starting with Junos OS Release 18.2R1, OpenConfig support through remote procedure call (RPC) and JTI is extended to support additional ON_CHANGE sensors for some endpoints under resource paths **/interfaces/interface/state** and **/interfaces/interface/subinterfaces/subinterface/state/**.

  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.

To provision a sensor to export data through 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.

To enable ON_CHANGE support, configure the sample frequency in the subscription as zero. When you create a subscription using a top-level container as the resource path (for example, **/interface**), leaf devices under the resource path **/interface** with ON_CHANGE support are automatically streamed based on events. Other leaf devices will not be streamed.

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.

[See Understanding OpenConfig and gRPC on Junos Telemetry Interface and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **J-Insight Device Monitor (MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, and vMX)**–J-Insight is a data-driven device monitoring solution that provides visibility and insight into the health of a running system. Starting with Junos OS Release 18.2R1, the J-Insight framework facilitates real-time monitoring of system resources for FPC FRUs. It also has been integrated with the existing connectivity error management infrastructure to normalize error detection, monitoring, and reporting. J-Insight is an on-premise system application that uses the Junos Telemetry Interface to continuously collect data that is reflective of the current state and health of the device component being monitored.

  [See J-Insight Device Monitor Overview.]

- **Service set and sessions support for Junos Telemetry Interface (JTI) (MX Series with MS-MICs and MS-MPCs)**—Starting with Junos OS Release 18.2R1, you can export service set and sessions statistics. These sensors provide visibility for IPsec services on different service complexes and nodes.

  Exported data is defined using an IP address and a UDP port. When an export interval expires, the most recent statistics collected by the sensors are gathered, placed in the payload of a UDP packet, and forwarded to a collector. A timestamp indicating when counters are read is included with the exported data to allow collectors to collate data. The timestamp also can determine if and when an event happened, such as a PIC hardware restart or if counters were cleared by means of the CLI.

  The following paths are supported:

  - **/junos/services/spu/servicesets/**
  - **/junos/services/spu/sessions/**

  For streaming statistics through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level.

  [See sensor and Configuring a Junos Telemetry Interface Sensor (CLI Procedure).]

Layer 2 Features

- **Support for Layer 2 and Layer 3 features (MX10008)—**Starting in Junos OS Release 18.2R1, MX10008 routers support the following Layer 2 and Layer 3 features:

- Layer 2 protocols including Layer 2 Ethernet OAM and virtual private LAN service (VPLS)
- Integrated routing and bridging (IRB)
- Multichassis link aggregation groups (MC-LAGs)
- Layer 3 routing protocols and MPLS
- Inline BFD
- Multicast

[See Layer 2 and Layer 3 Features on MX Series Routers]

Layer 3 Features

- **Multipoint support for ATM MIC with SFP (MX Series routers with MPCs and ATM MIC with SFP)**—Starting in Junos OS Release 18.2R1, MX Series routers with an ATM MIC (model number MIC-3D-8OC3-2OC12-ATM) with SFP can communicate with multiple devices through ATM links. With this multipoint support feature, ATM MIC can communicate with multiple Layer 3 peers in the ATM network. In earlier Junos OS releases, the ATM MIC communicates only with one Layer 3 peer.

  On an ATM MIC, the following configurations are required for multipoint support:

  - Configure the **multipoint** option at the **[edit interfaces *interface-name* unit *logical-unit-number*]** hierarchy level to communicate with multiple Layer 3 peers on ATM interface.
  - Configure the **multipoint-destination** option with its corresponding **vci** at the **[edit interfaces *interface-name* unit *logical-unit-number* family *family* address *address*]** hierarchy to enable multipoint support on ATM interface.

    The **Inverse-arp** configuration option is an optional configuration to enable inverse ARP for **multipoint-destination** at the **[edit interfaces at-fpc/pic/port unit logical-unit-number family family address address multipoint-destination address]** hierarchy level. Only responding to inverse ARP request is supported. Generation of Inverse ARP is not supported.

  [See Configuring a Point-to-Multipoint Connection on ATM MICs.]

MPLS

- **Interoperability of segment routing with LDP (MX Series)**—In an LDP network with gradual deployment of segment routing, some devices may not support segment routing, which can cause interoperability issues in the network. Starting in Junos OS Release 18.2R1, you can use OSPF or ISIS to enable segment routing devices to operate with the LDP devices that are not segment routing capable.

  To implement this feature using OSPF, an extended prefix link-state advertisement (LSA) with Range type, length, and value (TLV) for all the LDP prefixes is generated, and mapping routes corresponding to the prefix is installed in the inet.3 and mpls.0 routing tables.

  To implement this feature using ISIS, a server-client configuration is required under protocols ISIS and LDP, respectively, and routes from the inet.3 or inet.0 routing tables are used for stitching of segment routing LSP with an LDP LSP and vice-versa.

[See LDP Mapping Server for Interoperability of Segment Routing with LDP Overview.]

- **Support for reporting binding SIDs to a PCE (MX Series)**—Static non-colored segment routing LSPs have binding segment identifiers (SIDs) that are used for stitching multiple non-colored segment routing LSPs. Junos OS supports a maximum of five next hops for provisioning such segment routing LSPs.

  Starting in Junos OS Release 18.2R1, statically configured non-colored segment routing LSPs on the ingress device are reported to the Path Computation Element (PCE) through a Path Computation Element Protocol (PCEP) session. These non-colored segment routing LSPs may have binding SID labels associated with them.

  With this feature, the PCE can use this binding SID label in the label stack to provision PCE-initiated segment routing LSP paths.

  [See Static Segment Routing Label Switched Path.]

Network Management and Monitoring

- **SNMP MIB and trap support for group VPN members (MX Series)**—Starting in Junos OS Release 18.2R1, an SNMP MIB (jnxGdoiMIB) is added under jnxMibs, which is based on the GDOI MIB draft, draft-kamarthy-gdoi-mib-01. This MIB provides the SNMP MIB tables and notifications required for group VPN members. SNMP **get**, **get_next**, and **walk** functionality is added to the following tables:

  - jnxGdoiGroupTable
  - jnxGdoiGmTable
  - jnxGdoiGmKekTable
  - jnxGdoiGmTekSelectorTable
  - jnxGdoiGmTekPolicyTable

  Also, SNMP trap notifications are provided for the following events:

  - jnxGdoiGmRegister
  - jnxGdoiGmRegistrationComplete
  - jnxGdoiGmReRegister
  - jnxGdoiGmRekeyReceived
  - jnxGdoiGmRekeyFailure

  [See MIB Explorer and Standard SNMP MIBs Supported by Junos OS]

- **RPM timestamping extension on JNP10K-LC2101 (MX10008)**—Starting in Junos OS Release 18.2R1, JNP10K-LC2101 supports timestamping of RPM probes in the Packet Forwarding Engine host processer. You can enable this feature by including the **hardware-timestamp** statement at the [edit services rpm probe probe-name test test-name] hierarchy level.

  [See hardware-timestamp].

- **Support for RPM probes with IPv6 sources and destinations on JNP10K-LC2101 (MX10008)**—Starting in Junos OS Release 18.2R1, the RPM client router (the device that originates the RPM probes) can send probe packets to the RPM probe server (the device that receives the RPM probes) that contains an IPv6 address. To specify the destination IPv6 address used for the probes, include the **target (url *ipv6-url*|address *ipv6-address*)** statement at the **[edit services rpm probe owner test test-name]** hierarchy level. You can also define the RPM client or the source that sends RPM probes to contain an IPv6 address. To specify the IPv6 protocol-related settings and the source IPv6 address of the client from which the RPM probes are sent, include the **inet6-options source-address *ipv6-address*** statement at the **[edit services rpm probe owner test *test-name*]** hierarchy level.

- **Support for specifying a maximum hop count for RPM and TWAMP probes (MX Series routers)**—Starting in Junos OS Release 18.2R1, you can set a maximum hop count (TTL) for real-time performance monitoring (RPM) probes (both IPv4 and IPv6). This can be useful, for example, to restrict the scope of a given RPM probe so it cannot unintentionally monitor an alternative path to the destination, such as may occur following a BGP rerouting. Probes that exceed the number set for TTL are discarded.

  This TTL configuration is supported on Routing Engine-based RPM, MS-MPC based RPM, MS-MIC-based RPM, and Two-Way Active Management Protocol (TWAMP).

  [See RPM Overview and TTL (RPM probe).]

Restoration Procedures and Failure Handling

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (MX Series)**—Starting in Junos OS Release 18.2R1, 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 the 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

- **Topology-independent loop-free alternate for OSPF (MX Series)**—Starting in Junos OS Release 18.2R1, topology-independent loop-free alternate (TI-LFA) with segment routing provides MPLS fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. You can enable TI-LFA for OSPF by configuring the **use-post-convergence-lfa** statement at the **[edit protocols ospf backup-spf-options]** hierarchy level. When used with OSPF, TI-LFA provides protection against link failure and node failure.

  You can enable the creation of post-convergence backup paths for a given interface by configuring the **post-convergence-lfa** statement at the **[edit protocols ospf interface *interface-name* level *level*]** hierarchy level. The **post-convergence-lfa** statement enables link-protection mode.

You can enable **node-protection** mode for a given interface at the **[edit protocols ospf area *area* interface *interface-name* post-convergence-lfa]** hierarchy level.

[See Topology-Independent Loop-Free Alternate with Segment Routing for OSPF..]

- **Support for route leaking in OSPF stub areas (MX Series)**—Starting with Junos OS Release 18.2R1, route leaking is supported in OSPF when the router is overloaded, which allows redistribution of the external prefixes.

  Prior to Junos OS Release 18.2 , the external prefixes are not redistributed when OSPF is overloaded.

  You can now configure the following when OSPF is overloaded.

  - **allow-route-leaking** at the **[edit protocols <ospf | ospf3> overload]** hierarchy level to advertise the external prefixes with maximum cost.
  - **stub-network** at the **[edit protocols ospf overload]** hierarchy level to advertise stub network with maximum metric.
  - **intra-area-prefix** at the **[edit protocols ospf3 overload]** hierarchy level to advertise intra-area prefix with maximum metric.
  - **as-external** at the **[edit protocols <ospf | ospf3> overload]** hierarchy level to advertise external prefix with maximum metric.

  [See Understanding OSPF Overload Function.]

Services Applications

- **Traffic Load Balancer enhancements (MX Series with MS-MPCs)**—Starting in Junos OS Release 18.2R1, the Traffic Load Balancer (TLB) application supports the following enhancements:
  - 2000 TLB instances for virtual services that use the direct-server-return or translated mode
  - Tracing at the instance level or at the virtual services level
  - Display of real server up and down counts

  [See Configuring TLB.]

- **Port Control Protocol support for DS-Lite on MS-MPCs and MS-MICs (MX Series routers)**—Starting in Junos OS Release 18.2R1, Port Control Protocol (PCP) on the MS-MPC and MS-MIC supports DS-Lite. PCP provides a mechanism to control the forwarding of incoming packets by upstream devices such as NAT44 and firewall devices, and a mechanism to reduce application keepalive traffic.

  [See Port Control Protocol Overview.]

- **IKE and IPsec enhancements on MS-MPCs and MS-MICs (MX Series routers)**—Starting in Junos OS Release 18.2R1, the following enhancements are supported on MS-MPCs and MS-MICs:
  - You can configure the MX Series router to act only as an IKE responder. In this responder-only mode, the MX Series router does not initiate IKE negotiations, it only responds to IKE negotiations initiated by the peer gateway.

- You can configure the MX Series router to send only the end-entity certificate for certificate-based IKE authentication instead of the full certificate chain. This avoids IKE fragmentation.
- You can display the total elapsed time for a tunnel across security association rekeys (**Total uptime**) and the configured hard lifetime for a security association (**SA lifetime**) by running the **show services ipsec-vpn ipsec security-associations detail** command.

[See Configuring IKE Activation Time, Configuring IPsec Service Sets, and show services ipsec-vpn ipsec security-associations.]

- **Support for additional DS-Lite features on MS-MPCs and MS-MICs (MX Series routers)**—Starting in Junos OS Release 18.2R1, dual-stack lite (DS-Lite) running on MS-MPCs and MS-MICs adds support for the following features:
  - SIP ALG
  - Subscriber session limitation per subnet
  - DS-Lite service sets on AMS interfaces

  [See DS-Lite Subnet Limitation.]

- **Support of IPv6 probes for optimized CLI configuration of RPM tests (MX Series)**—Starting in Junos OS Release 18.2R1, you can also optimize the CLI configuration for RPM tests with IPv6 probes. Prior to Junos OS Release 18.2R1, you could only optimize the CLI configuration for RPM tests with IPv4 probes. Enter the **rpm-scale** statement at the **[edit services rpm probe *probe-owner* test *test-name*]** hierarchy level to generate multiple tests from a single configuration.

  [See Configuring RPM Probes.]

- **Inline JFlow support for EVPN traffic (MX Series)**— Starting in Junos OS Release 18.2R1, inline jflow supports sampling under the bridge family. Inline Jflow monitors traffic hitting the bridge family and reports the necessary fields in either version 9 or IPFIX format.

  A new family **bridge** is introduced under the **forwarding-options sampling instance** hierarchy that monitors all traffic hitting the VPLS or bridge family.

  [See Understanding Inline Active Flow Monitoring.]

Software Installation and Upgrade

- **ZTP support is added for MX VM host platforms (MX Series)**—In Junos OS Release 18.2R1, ZTP, which automates the provisioning of the device configuration and software image with minimal manual intervention, is supported on MX Series VM hosts. When you physically connect a supported device to the network and boot it with a factory configuration, the device attempts to upgrade the Junos OS software image automatically and autoinstall a configuration provided on the DHCP server.

  [See Zero Touch Provisioning.]

- **Unified ISSU support (MX10003)**—Starting in Junos OS Release 18.2R1, MX10003 routers support unified in-service software upgrade (ISSU). Unified ISSU enables you to upgrade from a particular Junos OS release to a later one with no disruption on the control plane and with minimal disruption of traffic.

MX10003 supports unified ISSU upgrade of the complete package including VMHOST Linux. Use the command **request vmhost software in-service-upgrade** to perform a unified ISSU upgrade. [See Understanding the Unified ISSU Process]

NOTE:

- Starting in Junos OS Releases 18.2R1, MX10003 supports unified ISSU.
- MX10003 does not support upgrading only the Junos OS image using the **request system software in-service-upgrade** command.
- Unified ISSU is not supported on MACsec MIC (JNP-MIC1-MACSEC).
- Unified ISSU is not supported for the interfaces that are configured with 1-Gigabit Ethernet mode.
- Unified ISSU is not supported on timing protocols (for example, Precision Time Protocol and Synchronous Ethernet), MACsec protocols, and BBE protocols.
- The **MAC statistics** (retrieved using the show interfaces extensive command) are reset during unified ISSU which means that the **MAC statistics** command does not provide the correct statistics after unified ISSU.

Subscriber Management and Services

- **Local authentication and authorization for subscribers (MX Series)**—Starting in Junos OS Release 18.2R1, you can enable local authentication and limited local authorization for individual subscribers instead of using external authentication and authorization servers. To enable local authentication, specify the **password** option of the **authentication-order** statement in the access profile. Define the local password for the subscriber with the **password** option of the **subscriber *username*** statement in the access profile. You can configure up to 100 subscribers for local authentication chassis-wide. You can optionally configure local authorization with other options of the **subscriber *username*** statement. Local authentication statistics are displayed by the **show network-access aaa statistics authentication detail** and **show network-access requests statistics** commands.

  [See Configuring Local Authentication and Authorization for Subscribers.]

- **Nonterminating filter actions next-ip and next-ip6 supported in dynamic profiles (MX Series)**—Starting in Junos OS Release 18.2R1, the firewall filter actions **next-ip** and **next-ip6** are available in dynamic profiles. Already supported for static profiles, these nonterminating actions direct packets matching a given filter to the specified IPv4 or IPv6 destination address. When the filters including these actions are in dynamic service profiles, you can create user-defined variables to parameterize the associated address and the optional routing instance name.

  [See Parameterized Filter Nonterminating and Terminating Actions and Modifiers.]

- **Automatic validation of DHCPv6 client MAC addresses to reduce session hijacking (MX Series)**—Starting in Junos OS Release 18.2R1, the DHCPv6 local server and relay agent automatically attempt to validate

a client's MAC address to prevent accepting packets from malicious clients that attempt to hijack the client session.

When DHCPv6 local servers and relay agents receive a solicit message from a client to establish a session, they extract the client MAC address (link-layer address) from the message and add it to a local table that maps MAC addresses to client IPv6 addresses or prefixes. They use this table to compare MAC addresses received in subsequent messages from the client to validate whether the client is known; if not, it is assumed to be malicious and the control packet is dropped. Because the packet has failed MAC validation, the client MAC validation counter is incremented.

[See DHCPv6 Client MAC Address Validation to Prevent Session Hijacking.]

- **DHCP short-cycle protection to reduce excess loading (MX Series)**—Starting in Junos OS Release 18.2R1, you can enable the router to identify DHCP clients that have short logins or continually fail to connect; the router then drops subsequent requests from these clients until a lockout timer expires. For users that repeatedly log in frequently and briefly, the initial lockout time is short enough to have no noticeable impact. As these brief logins continue, the lockout time is exponentially increased.

  [See DHCP Short Cycle Protection Against Frequent Brief or Failed Client Sessions.]

- **Support for direct PCC rule activation by a PCRF (MX Series with MS-MPCs)**—Starting in Junos OS Release 18.2R1, a policy and charging rules function (PCRF) server can directly activate a policy and charging control (PCC) rule that is configured on the MX Series router. To activate a PCC rule, the PCRF sends a Rule-Install-Name AVP over the Gx interface to the MX Series router. PCC rules define the treatment to apply to subscriber traffic (for example, setting the maximum bit rate) based on the application being used by the subscriber (for example, Facebook) or based on the Layer 3 and Layer 4 service data flow information for the IP flow (for example, the source and destination IP addresses).

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

- **Dynamic profile enhancements to support migration of static, terminated IPv4 PPP subscribers (MX Series)**—Starting in Junos OS Release 18.2R1, the following enhancements support dynamic profiles for static subscribers:

  - CPE-sourced subscriber address—You can direct jpppd to use the IP address supplied by the client in an incoming IPCP configure-request message rather than assigning another address by configuring your RADIUS server to Framed-IP-Address attribute (8) with the wildcard value of 255.255.255.255.
  - Tag2 for static routes—For PPP subscribers that use static routes with a tag2 attribute for MP-BGP, you can configure your RADIUS server to include the tag2 attribute in the Framed-Route attribute [22] when it authenticates a subscriber. Alternatively, you can configure the dynamic profile to provide a specific tag2 value for a specific access route prefix.
  - Local authentication—For clients that do not support authentication protocols such as PAP and CHAP, you can configure usernames and passwords locally. You can define the name based on one or more of the following: MAC address, agent circuit identifier, agent remote identifier, and domain name. The router uses these values when it contacts the RADIUS server for authentication.

  [See Migrating Static PPP Subscriber Configurations to Dynamic Profiles Overview.]

- **Support for multipoint LDP to utilize distributed IGMP to signal an MPLS-based core has been added to Enhanced Subscriber Management (single-chassis MX Series routers with MPC2E, MPC3E, MPC5E, or MPC7E)**—Starting in Junos OS Release 18.2R1, support for multipoint LDP inband signaling to interwork with distributed IGMP has been added. As such, two separate PIM domains can be interconnected by an MPLS-based core (that is, a PIM-free core). One application of multipoint LDP inband signalling is to carry IPTV multicast traffic on an MPLS backbone.

  To enable the interworking, **chassis network-services enhanced-ip** must first be configured. Then you need to set the **igmp** or **mld** interface for **distributed**, and enable **mldp-inband-signalling** at the PIM hierarchy so PIM acts as a multipoint LDP inband edge router:

  ```
  [edit dynamic-profiles profile-name protocols igmp|mld interface layer 3 interface name distributed]
  [edit protocols pim mldp-inband-signalling]
  ```

  You can the run the **show pim source** command to confirm that distributed multipoint LDP is working (look for **Upstream neighbor via MLDP-inband**).

  [See Understanding Distributed IGMP and Enhanced Subscriber Management Overview]

- **BNG support for cascading DSLAM deployments over bonded DSL channels (MX Series)**— Starting in Junos OS Release 18.2R1, Passive Optical Network (PON) access technologies are supported with four levels of quality-of-service (QoS) scheduler hierarchy for residential subscribers in a BBE deployment. This feature extends the Access Node Control Protocol (ANCP) implementation to handle network configuration for residential customers that use PON as the broadband access technology for both CuTTB and FTTB. ANCP uses a statically controlled traffic-control profile on the interface-set for shaping at the subscriber level at the intermediate node to which the subscribers are connected. New DSL types are provided to support access line rate adjustment for the new access technologies.

  [See Understanding BNG Support for Cascading DSLAM Deployments Over Bonded DSL Channels

  [See access-line (Access Line Rate Adjustment).]

  A new RADIUS VSA, **inner-tag-protocol-id** 26-211 is introduced to fetch the inner VLAN Tag Protocol Identifier value for L2BSA subscribers to enable maintaining one dynamic profile instead of two separate dynamic profiles. A new Junos OS dynamic profile variable *$junos-inner-vlan-tag-protocol-id* allows a VLAN map's **inner-tag-protocol-id** to be set by RADIUS or a predefined default value provided in the configuration.

  [See Junos OS Predefined Variables That Correspond to RADIUS Attributes and VSAs.]

- **Global range setting for initial router advertisement intervals (MX Series)**—Starting in Junos OS Release 18.2R1, you can configure override options to set a global range from which the router randomly selects an interval for each interface for only the initial three router advertisements that the router sends when the router becomes available on that interface. This enables you to set a range that results in a very short interval for these advertisements without affecting subsequent advertisements set by the router.

  In earlier releases, you can configure an interval range only per interface and the range settings apply to all router advertisements that the router sends.

[See Configuring an Interval Range for Unsolicited Router Advertisements to IPv6 Neighbors.]

System Management

- **New tool to detect high CPU utilization for routing protocol process (MX Series)**—Starting in Junos OS Release 18.2R1, a flight recorder tool is introduced to gather historical data on when the CPU utilization for routing protocol process on a device was high and what processes caused the high utilization. The tool collects snapshots of data, enabling detection of high CPU usage and faster resolution of issues.

  Because some of the high CPU utilization cases are intentional or expected, you can enable and disable the flight recorder tool to avoid false alarms.

  [See request flight-recorder set high-cpu and show flight-recorder status.]

User Interface and Configuration

- **Support for displaying ephemeral configuration data with filtering (MX Series)**—Starting in Junos OS Release 18.2R1, the **show ephemeral-configuration** command enables you to specify the scope of the configuration data to display. To filter the displayed configuration data, append the statement path of the requested hierarchy to the command.

  [See Displaying Ephemeral Configuration Data in the Junos OS CLI.]

VPN

- **Increased number of supported routing instances (MX 960 and MX 2020)**—Starting in Junos OS Release 18.2R1, Junos OS supports up to 16,000 VPLS routing instances with 128,000 (FEC 128) hierarchical VPLS pseudowires.

  [See Configuring VPLS Routing Instances.]

SEE ALSO

Changes in Behavior and Syntax | 134

Known Behavior | 145

Known Issues | 153

Resolved Issues | 176

Documentation Updates | 222

Migration, Upgrade, and Downgrade Instructions | 223

Product Compatibility | 230

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 18.2R3 for MX Series..

Class of Service (CoS)

- **Junos OS 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

- **Change in the output for show evpn instance and show evpn database**—Starting in Junos OS Release 18.2R1, the output for a**show evpn instance** and **show evpn database** displays a local interface with an interface name of .local..*number*. and no configuration. This interface is created to support configuration fault management (CFM). For example, **show evpn instance** displays the following sample output:

```
Number of local interfaces: 2 (2 up)
    Interface name  ESI                             Mode             Status
AC-Role
    .local..9       00:00:00:00:00:00:00:00:00:00  single-homed     Up
Root
```

- **Support for a VNI of zero**—Starting in Junos OS Release 18.2R3, Junos OS 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 18.2R3, 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

- **No error codes displayed for PFE errors (MX Series)**—Starting in Junos OS Release 18.2R1, on MX Series routers, the **show chassis alarms** output does not display error codes for Packet Forwarding Engine -related errors. You can use the following commands to view more details of the errors that caused the alarms:
  - **show chassis errors active**
  - **show chassis errors active detail**

- **User confirmation prompt for configuring the sub-options of request vmhost commands (MX Series)**—While configuring the following **request vmhost** commands, the CLI now prompts you to confirm a **[yes,no]** for the sub-options also.
  - **request vmhost reboot**
  - **request vmhost poweroff**
  - **request vmhost halt**

  In previous releases, the confirmation prompt was available for only the main options.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (MX Series)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

Infrastructure

- **Change in support for interface-transmit-statistics statement (MX Series)**—You cannot configure aggregated Ethernet interfaces to capture and report the actual transmitted load statistics by using the **interface-transmit-statistics** statement. Aggregated Ethernet interfaces do not support reporting of the transmitted load statistics. The **interface-transmit-statistics** statement is not supported in the aggregated Ethernet interfaces hierarchy. In earlier releases, the **interface-transmit-statistics** statement was available in the aggregated Ethernet interfaces hierarchy but not supported.

Interfaces and Chassis

- 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 18.2R3, integrated routing and bridging (IRB) is not supported on a pseudowire

subscriber (PS) logical interface. Hence, you cannot add IRB to the bridge domain with a PS interface, that is, you cannot configure IRB and a PS interface in the same bridge domain.

Adding IRB to a bridge domain with a 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)**—Starting in Junos OS Release 18.2R3, 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)**—Starting in Junos OS Release 18.2R3, 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 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

- **Junos XML protocol <open-configuration> operation no longer emits an uncommitted changes warning (MX Series)**—Starting in Junos OS Release 18.2R1, the Junos XML protocol **<open-configuration>** operation does not emit an "**uncommitted changes will be discarded on exit**" warning message when opening a private copy of the candidate configuration. However, Junos OS still discards the uncommitted changes upon closing the private copy.
- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (MX Series)**—Starting in Junos OS Release 18.2R2, 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.

Junos Telemetry Interface

- **Change to the configuration location for gRPC-based sensor subscriptions from an external collector (MX Series)**—Starting in Junos OS Release 18.2R1, when an external streaming server, or collector, provisions sensors to export data through gRPC on devices running Junos OS, the sensor configuration is committed to the **junos-analytics** instance of the ephemeral configuration database, and the configuration can be viewed by using the **show ephemeral-configuration instance junos-analytics** operational command. In earlier releases, the sensor configuration is committed to the default instance of the ephemeral configuration database.

MPLS

- **Display of Route Session-ID count in the show rsvp version command output**—Starting in Junos OS Release 18.2R1, the **show rsvp version** command output displays the **Route Session-ID count** output field by default, even when there are no session IDs associated with the RSVP ingress routes. In such cases, the **Route Session-ID count** value is zero (0).
- **Support for inet.0 and inet.3 labeled unicast BGP route for protocol LDP (MX Series)**— Starting in Junos OS Release 18.2R1, LDP egress policy is supported on both inet.0 and inet.3 routing Information bases (RIBs) also known as routing table for labeled unicast BGP routes. If a routing policy is configured with a specific (inet.0 and inet.3) RIB, the egress policy is applied on the specified RIB. If no RIB is specified and a prefix is present on both inet.0 and inet.3 RIBs for labeled unicast BGP routes, then inet.3 RIB is preferred. However, prior to Junos OS Release 12.3R1 and starting with Junos OS Release 16.1R1, LDP egress policy is always preferred on inet.0 RIB and support for inet.3 RIB egress policy for labeled unicast BGP routes was disabled. In Junos OS Release 12.3R1 and later releases up to Junos Release 16.1R1, LDP egress policy was supported in inet.3 RIBs, in addition to inet.0 RIBs, for labeled-unicast BGP routes.
- **New option in show mpls lsp autobandwidth command**—In Junos OS Release 18.2, a new option—**name *lsp-name***—is introduced in the **show mpls lsp autobandwidth** command to specify the name of the LSP for which the autobandwidth information is displayed. With the **name** option, the autobandwidth information specific to the LSP name that has been provided can be obtained in the command output.

  [See show mpls lsp autobandwidth.]
- **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.

- 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 18.2R2, with a similar configuration, the **HelloInterval** output field displays 0 as the hello interval.

- 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 18.2R2, 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.

- **New debug statistics counter (MX Series)**—The **show system statistics mpls** command has a new output field, called **Packets dropped, over p2mp composite nexthop**, to record the packet drops over composite point-to-multipoint next hops.

Network Management and Monitoring

- Starting in Junos OS Release 18.2R1, there must be no space in the password for configuring the Network Time Protocol (NTP) authentication-key. For example **user@host# set system ntp authentication-key 10 type md5 value "ABCDjuniper"**.

  Prior to Junos OS Release 18.2R1, the NTP authentication or password was successfully configured with a space added in the password. For example **user@host# set system ntp authentication-key 10 type md5 value "ABCD juniper"**.

- **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 (MX Series)**—In Junos OS Release 18.2R1, 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.]

- A decrease in the MPLS label-switched path (LSP) statistics pauses the SNMP MIB **mplsLspInfoAggrOctets** count for one MPLS statistics gathering interval. In such cases, the **mplsLspInfoAggrOctets** value is updated only after completing one more interval of the MPLS statistics gathering.
- **Junos OS does not support management of YANG packages in configuration mode (MX Series)**—Starting in Junos OS Release 18.2R2, 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 18.2R2, 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 18.2R2, 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.]

Platform and Infrastructure

- **NTP Boot Server configuration (MX204, MX960, MX10003, MX10002, MX10016, MX10000, MX480, MX104, MX10008, MX240, MX2010, MXTSR80, MX80, MX2008, MX150, and MX2020)**— Use **set ntp server *address*** command to set the correct time when we boot the router instead of boot-server *address*.

  [See Synchronizing and Coordinating Time Distribution Using NTP.]

Routing Protocols

- **Change in the default behavior of advertise-from-main-vpn-tables configuration statement**—BGP now advertises EVPN routes from the main bgp.evpn .0 table. You can no longer configure BGP to advertise the EVPN routes from the routing instance table. In earlier Junos OS Releases, BGP advertised EVPN routes from the routing instance table by default.

  [See advertise-from-main-vpn-tables.]

- **Modified output of show route forwarding-table**—Starting in Junos OS Release 18.2R1, the output of **show route forwarding-table** command does not display the next-hop address for static routes that use point-to-point (P2P) interfaces.

  [See show route forwarding-table.]

- **MPLS configuration mandatory for indirect next-hop interfaces**—It is mandatory for an indirect next-hop's forwarding interface to have family MPLS configured. In a BGP network if the MPLS configuration for an indirect next-hop's forwarding interface is deleted or when the BGP labeled unicast interface is deactivated, all routes with indirect next hop undergo a route resolution again, which might impact traffic routing until the route resolution is completed. In earlier Junos OS releases when family MPLS was deleted, the indirect next-hop route was removed from the forwarding table and could not be recovered even when MPLS was reactivated.

Services Applications

- **Change in error message displayed while fragmenting or defragmenting IPv6 GRE tunnel interface (MX Series routers)**—Starting in Junos OS Release 18.2R3, 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)**—Starting in Junos OS Release 18.2R3, you can send host-generated traffic on a GRE over GRE tunnel. However, when the path

maximum transmission unit (PMTU) is updated for the outer GRE tunnel, MTU for the inner GRE tunnel is not corrected.

- **New syslog message displayed during NAT port allocation error (MX Series routers with MS-MPC)**—With address pooling paired (APP) enabled, an internal host is mapped to a particular NAT pool address. In case all the ports under a NAT pool address are exhausted, further port allocation requests from the internal host results in a port allocation failure. The following new syslog message is displayed during such conditions:

  JSERVICES_NAT_OUTOF_PORTS_APP

  This syslog message is generated only once per each NAT pool address.

Software Defined Networking

- **Installation or upgrade using remotely located installation package (MX480, MX960, MX2010, MX2020, MX2008)**—While performing Junos OS installation or upgrade on the base system (BSYS) or guest network function, if you provide a URL to the remotely located installation package (for example, an ftp file) in the command **request system software add *package-file-path***, the router locally copies the package, performs checks such as multi-version compatibility checks on the package, and then installs the package. The installation process is aborted if any errors are found during the checks. Previously, if you tried to perform installation or upgrade using a remotely located file, the router would skip multi-version checks and display an error message, but would not abort the installation process.

  [See Junos Node Slicing Upgrade]

Software Installation and Upgrade

- **New DHCP option introduced for ZTP retry (MX Series)**—Starting in Junos OS Release 18.2R1, a new DHCP option is introduced to set the timeout value for the file downloads over FTP. If the **transfer-mode** is set as FTP, the default value for the timeout is automatically set as 120 minutes. That is, if the FTP session gets interrupted due to loss of connectivity in the middle of a file transfer, it will time out after 120 minutes and ZTP will attempt to retry the file-fetching process. This value can be overridden using the DHCP option as follows:

  ```
  option NEW_OP.ftp-timeout code 7 = text;
  option NEW_OP.ftp-timeout “val”;
  ```

  where **“val”** is the user configurable timeout value in seconds and must be provided within quotes (for example,"val”).

- **ZTP is supported on MX PPC platforms (MX Series)**—As of Junos OS Release 18.2R1, 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

- **Wildcard supported for show subscribers agent-circuit-identifier command (MX Series)**—Starting in Junos OS Release 18.2R1, you can specify either the complete ACI string or a substring when you issue the **show subscribers agent-circuit-identifier** command. To specify a substring, you must enter characters that form the beginning of the string, followed by an asterisk (*) as a wildcard to substitute for the remainder of the string. The wildcard can be used only at the end of the specified substring.

  [See show subscribers.]

- **Changed behavior for framed routes without a subnet mask (MX Series)**—Starting in Junos OS Release 18.2R1, the router connects the session but ignores a framed route when it is received from RADIUS in the Framed-Route attribute (22) without a subnet mask.

  In earlier releases, the router installs the framed route with a Class A, B, or C subnet mask depending on the value of the first octet. When the octet < 128, the mask is /8; when 128 <= octet < 192, the mask is /16; and when the octet >= 192, the mask is 24.

- **DHCPv6 lease renewal for separate IA renew requests (MX Series)**—Starting in Junos OS Release 18.2R2, 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.]

- **Bandwidth options match for inline services and tunnel services (MX Series)**—Starting in Junos OS Release 18.2R2, 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)].

- **Disabling a pseudowire underlying interface (MX Series)**—Starting in Junos OS Release 18.2R2, 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.]

- **ICMP error message rate limit increased (MX Series)**—Starting in Junos OS Release 18.2R3, 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.

- **Out-of-address SNMP trap requires thresholds to be configured (MX Series)**—Starting in Junos OS Release 18.2R3, 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.

User Interface and Configuration

- **Changes to the show ephemeral-configuration command (MX Series)**—Starting in Junos OS Release 18.2R1, the **show ephemeral-configuration** operational mode command has the following changes:

  - To display the configuration data in the default instance of the ephemeral configuration database, issue the **show ephemeral-configuration instance default** command. In earlier releases, ephemeral configuration data for the default instance is displayed using the **show ephemeral-configuration** command.

  - To display the configuration data in a user-defined instance of the ephemeral configuration database, issue the **show ephemeral-configuration instance *instance-name*** command. In earlier releases, ephemeral configuration data for a user-defined instance is displayed using the **show ephemeral-configuration *instance-name*** command.

- To view the complete post-inheritance configuration merged with the configuration data in all instances of the ephemeral database, issue the **show ephemeral-configuration merge** command. In earlier releases, the merged view is displayed using the **show ephemeral-configuration | display merge** command.

- **Change to the maximum number of user-defined instances supported by the ephemeral configuration database (MX Series)**—Starting in Junos OS Release 18.2R1, devices running Junos OS that support configuring the ephemeral configuration database enable configuring a maximum of seven user-defined instances of the ephemeral database. In earlier releases, you can configure up to eight user-defined instances. User-defined instances are configured using the **instance *instance-name*** statement at the **[edit system configuration-database ephemeral]** hierarchy level.

SEE ALSO

New and Changed Features | **112**

Known Behavior | **145**

Known Issues | **153**

Resolved Issues | **176**

Documentation Updates | **222**

Migration, Upgrade, and Downgrade Instructions | **223**

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 18.2R3 for MX Series.

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

EVPN

- In scaled EVPN VPWS configurations (approximately 8000 EVPN VPWS), during a Routing Engine switchover, rpd scheduler slip messages might be seen. PR1225153

Forwarding and Sampling

- On an I-chip that for an IP->MPLS case, the PTYPE is carried over a fabric as an IP address, but the egress MPLS features such as filters are not executed. So traffic will not reach the hit MPLS filter and matches in Inet filter PR751618
- As per the investigation from RPD : we have is an interface for a direct route starting in ifdown condition. The remote side is then brought up, so I/F goes to ifup. Since it is a direct route, rpd does not install the route or nexthop. It receives that info from the kernel, and just updates a nexthop in rpd local storage. route and nexthop for the interface are taken care of in the kernel. There is no route change in rpd. route_record depends on route flash to find out about updates. 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 all 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. Workarounds: ++++++++++++ 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

- Pl. see AT PR1403182

General Routing

- Source-prefix filtering and protocol filtering of the carrier-grade NAT sessions are incorrect. For example, **show services sessions extensive protocol udp source-prefix <0:7000::2>** displays incorrect filtering of the sessions. PR1179922
- 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, next hop> pair that looks legal from the application standpoint, but KRT is not able to download it to the kernel, because the kernel rejected the next-hop, the <route, next-hop> sort of gets stuck in the rpd process. In the meantime, the previous version of the route (L-ISIS in this case) that was downloaded still lingers in the kernel and the Packet Forwarding Engine. PR1253778
- On Junos OS, the FPC might get stuck sometimes in offline state with the reason **Restarted by cli command** after restarting the FPC immediately after restarting chassisd. This issue is because of the fact that it takes some time for the system to stabilize after chassisd restarts. Though chassisd provides the FPC status and accepts the commands, it might take some time to initialize the process. Therefore, wait until all the PIC status are also available before issuing any command that makes FPC online, offline and restart. PR1275530
- 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 the GRE interface in case the bridge is configured with CFM. PR1275833
- At reboot, the RHEL 7.3 servers report **libvirtd[6282]: segfault at 10 ip 00007f87eab09bd0. No core file isleft** and no operational impact is known. PR1287808
- Support for enterprise profile is provided only for 10-Gigabit Ethernet interfaces. Using 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. PR1310048
- The InputInt field of MPLS-V4 data records reports the SNMP index value of the LSI interface instead of the ingress physical interface. PR1312047
- Sometimes the 1-Gigabit Ethernet interface might remain down after the following events:
  - Two are interfaces connected on a loopback on the 12-port QSFP28 TIC on an MX10003 or on a 4-Port QSPF28 fixed PIC on an MX204 and configured as 1-Gigabit Ethernet interfaces.
  - Two MX10003 devices are connected back-to-back and rebooted at the same time with 1-Gigabit Ethernet interfaces on a 12-port QSFP28 TIC. PR1312403

- When cmerror disables the Packet Forwarding Engine, it does not power off the EA and the 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
- Memory optimizations were done in carrier-grade NAT to increase the per session memory usage as in Junos OS Release 17.4R2 and the fix was committed through this PR. With the fix, scaling numbers are improved and 6M sessions are established. But the given fix is not enough to maintain 6M scaling sessions across other releases starting in Junos OS 18.1 Release. In these releases, PCP-related and other feature changes also have gone in. These changes have resulted in reduction in base memory and hence memory per session is decreased. Further per session memory utilization optimization in NAT is difficult through the PR. So we need collective effort from carrier-grade NAT services as well infra side to optimize it further. As per current scenario we cannot support more than 5.5 M sessions with APP/EIM/EIF enabled. PR1328510
- When a packet enters an FTI tunnel, copying the inner packet's TTL into the outer header implies that any subsequent packet drop inside the tunnel is conveyed to the source of the original packet. This is not handled in the Packet Forwarding Engine currently. So the inner packet TTL is not copied to the outer encapsulated packet header. PR1338467
- QSFP+-40G-CU3M is not supported on the MX10003 router. PR1341969
- Some vmhost commands are missing on Zero Touch Provisioning (ZTP) for MX Series platforms with VM host support (that is next-generation Routing Engines, such as RE-S-X6-64G). This might cause the ZTP for vmhost images is failed on this kind of platform. PR1343338
- The Routing Engine boots from the secondary disk in the following instances:
  - press the reset button, on the RCB front panel, while Routing Engine is booting up but before Junos OS is up
  - upgrade software, by booting from the network using the request vmhost reboot network command, and the system fails to boot from the network
  - upgrade BIOS and the upgrade fails d) reboot and the system hangs before Junos is up. PR1344342
- After disabling the laser for CWDM optics, optics diagnostics will not report o/p power low and laser current low alarm/warnings. PR1349258
- After GRES there is no IPDR in the flatfile.PR1386148
- IDS aggregate configuration will not be considered for the installation of the IDS dynamic filter. PR1395316
- When a member link gets deleted or deactivated from an aggregated Ethernet bundle or the link goes down on the inline BFD session that is currently established, the BFD session might flap. This is a day-one limitation of the inline BFD design on PTX. PR1401342

- Junos OS does not perform a VLAN-ID check at the egress; VLAN-ID check is only performed at the ingress. PR1403730
- When SSH keys are generated during downgrade or upgrade of an image (usually on the first boot), <output> XML tags are visible in the messages. Taking out the xml tags might cause issues in netconf session. This is a minor cosmetic issue, hence does not have impact on the functionality. PR1432464
- On MX10003 platform with no MSATA device, xSTP topology change can be seen during FRU upgrade state in ISSU. PR1435397

Infrastructure

- An MX80 router might restart when the system memory usage is high. This restart can manifest in two ways, in both cases the system will restart. Both cases are related to the USB storage device present on the router. 1) The kernel core file is generated after restart. An example is shown under "Example stack" 2) The watchdog restart is triggered and no kernel core file is generated. Console output as described under “Example console output" is seen. There will be a gap of approximately 12 minutes in the message log, during the watchdog timeout. PR921062

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
- 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
- In case of hw-assisted-pm mode of operation at the responder, it takes a few milliseconds or seconds (based on the programmed scale) to program inline-responder entries once CCM comes up. So until the inline-entry corresponding to an SLM session does not get programmed, a response will not be sent back to the originator and originator will see a loss. Once an inline-responder entry gets programmed, responses will be sent back to the originator. PR1311963
- At JDM install time, each JDM instance generates pseudorandom 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 pseudorandom MAC addresses to be used as the chassis MAC address pool for the forwarding interfaces of that GNF. After they are generated, JDM and GNF MAC addresses are persistent, and will be deleted only 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.

- In MX10008 routers, the fabric is referred to as either Switch Interface Board (SIB) or Switch Fabric Board (SFB). The **show chassis hardware** output uses both SIB and SFB to refer to the fabric. The outputs of the commands **show chassis sfb errors** and **show chassis alarms** use SIB to refer to the fabric.
- **Error thrown when router configuration updated on live system**—In Junos OS Release 18.2R1, 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.
- It was identified that LCP Echo-Replies with an invalid Identifier from Client are all accepted by BNG, and thereby from the outside it looks like the BNG is not completely confirming to the following statement of the PPP standard (RFC 1661) PR1413777

Platform and Infrastructure

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

Routing Protocols

- Continuous soft core files might be generated due to a **bgp-path-selection** code. The routing protocol process (rpd) forks a child and the child asserts to produce a core file. The problem is with route ordering and it is automatically corrected after collecting the soft-assert-core file, without any impact to the traffic or service. PR815146
- When a Junos OS aggregation gateway uses an IPv6 address as a next-hop for IPv4 aggregates announced to the downstream, it might attract traffic prematurely before Packet Forwarding Engines are programmed with more specific IPv4 routes. This happens when the IPv6 address is advertised in the BGP **inet6-labeled-unicast** family. PR1220235
- The rpd-Packet Forwarding Engine is out-of-sync during MoFRR convergence. PR1284463

- Degradation is seen in BGP v4/v6 Routing Engine delete time when compared with Junos OS Release 17.2R1. PR1289582
- BGP peer flapping is seen when a 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
- BGP input and output threading are added in Junos OS Release 16.1R1 whereby BGP writes were batched to improve efficiency. This might sometimes lead to some latency in sending BGP update while reacting to certain network events. PR1332301
- While performing manual switchover of Routing Engine in a GRES/NSR enabled switchover or automatic switchover during NSSU upgrade in QFX Series Virtual Chassis, occasionally, the rpd process on the old master Routing Engine and new backup Routing Engine might undergo termination (instead of quiet exit due to a change in master Routing Engine) causing flapping of BGP peers. PR1417595

Services Applications

- We recommend that you do not configure **ms- interface** when an AMS bundle in one-to-one mode has the same member interface. PR1209660
- Inconsistent content might be observed in the access line information between ICRQ and PPPoE messages. PR1404259
- 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 (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 Installation and Upgrade

- **Unified ISSU not supported with an active RPM configuration**—If you have an active real-time performance monitoring (RPM) configuration, you cannot perform a successful unified in-service software upgrade (ISSU) to a Junos OS Release 18.2. The warning **ISSU is not supported for RPM configuration** appears.

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.

User Interface and Configuration

- On all devices running Junos OS, under **set interfaces *interface-range***, if you try to expand a valid interface name with auto-completion (TAB), then it might result in following error: **error: invalid value: <interface name>**. PR1353741

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 for the MX 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)

- While configuring the **rate-limit-burst** statement in the CoS hierarchy, the commit needs to push an update for CoS code handling on all the Packet Forwarding Engines and during this time, if an interface settings (internal attributes for an interface) is found to be NULL. Interface settings are usually stored in a memory location and the pointer to it becomes NULL because cosd does not check for the NULL values and results in segmentation fault. Channelized interface setting is found to be NULL for channelized interfaces, but the CoS code handling the configuration statement **rate-limit-burst** in the Packet Forwarding Engine is de-referenced by the setting without performing a NULL check, resulting in generation of a core file. A fix is added only to de-reference if the pointer is not null and as a result the crash is avoided. PR1425667

EVPN

- The Layer 2 address learning process (l2-ald) might generate a core file in a scaled Layer 2 setup, including 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 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
- In EVPN-VXLAN scenario, ARP table information is not synchronized on two Spines after reconfiguring an end host on a multihomed CE interface from IP1/MAC1 to IP1/MAC2. PR1330663
- When VTEP scale of more than 200 is used in Junos OS Release 18.1R1, VTEPs might not come up for all the tunnels and might impact traffic. PR1342175
- On devices running Junos OS software, the l2ald daemon might crash during the MAC address processing. The MAC learning process will be impacted during the period of l2ald crash. The l2ald recovers itself. PR1347606
- Type 2 EVPN routes are missing after deactivating or activating the protocol EVPN. PR1362598
- The following error is observed when trying to execute **vxlan ping overlay** through the RPC command: **ping-overlay: illegal option X**. It works through CLI but not through RPC . PR1373025
- When EVPN is configured with CoS-based forwarding (CBF), traffic might be lost for the CBF services. PR1374211

- The statement **forwarding-options multicast-replication evpn** is available in Junos OS Releases 17.4, 18.2, 18.3, and later. Multicast in EVPN in centrally-routed modes is not supported in the releases where this statement is missing. PR1381030
- After a MAC move from a local interface to a remote MAC, the **show bridge show evpn statistics** command reports the wrong number of MACs learned on an interface. The **show bridge/evpn mac-table count** command provides the accurate number of MACs learnt. PR1432293

Forwarding and Sampling

- Heap memory leaks occur on the DPC when the flow specification route is changed. PR1305977
- Firewall filter is not applied as input filter to an extended port when used for Layer 2 VPN. PR1311013
- The **show firewall filter** command does not display policer counters for filters that reference three-color policers. PR1364673

General Routing

- There is a timing issue between the Junos OS software and the I2C controllers on an MPC5E during a reboot. The software has been corrected to wait for I2C controllers to be ready before the software starts monitoring the voltage levels and current levels. PR1051902
- In a scaled setup, remnant routes might be seen in the old master Routing Engine after a Routing Engine switchover in a non-GRES scenario because the rpd process in the old master Routing Engine might not have enough time to clean all the routes from the kernel. In this case, there will be a convergence delay (in minutes) when this backup Routing Engine becomes master again. The length of the delay depends on the number of routes (for example, a 6M routes environment with approximately 1M remnant routes might have a 6-minute delay). PR1075404
- If a Layer 3 interface is receiving a GRE encapsulated packet and the interface has two filters attached in ingress, **Family any** with action as mirror and **Family inet** with action as decapsulate GRE, then the expected behavior is that the mirrored copy must have the GRE headers as well. As a workaround, deactivate or disable the decapsulate GRE action of the filter. PR1090854
- SIP session fails when the IPv4 SIP client in the public network initiates a SIP call with the IPv6 SIP client in the private network. PR1139008
- SMID daemon stops responding to the management requests after a jl2tpd (L2TP daemon) crash on an MX960 BNG. PR1205546
- In a rare condition, multiple interrupts are not handled properly on an MX Series platform with MPC7E, MPC8E, or MPC9E, and a PTX platform with FPC3-PTX-U2/FPC3-PTX-U3, which might create a core file. This condition is difficult to reproduce. As a workaround, the interrupt code is optimized to avoid the unnecessary call to prevent the issue. PR1208536

- When an MPC is removed while the card is online, the link error column in the **show chassis fabric summary extended** command output shows YES for all fabric planes. Conversely, when an MPC is taken offline using the CLI command, the output shows correctly. PR1214611
- The following error messages occur during GRES and unified ISSU: **syslog errors @ agentd_rts_async_rtbm_msg : FLM : Failed to create private**. PR1232636
- The MX104 Routing Engine might be stuck in boot loop after disabling the interface fxp0 in configuration. PR1253155
- Load balancing is uneven across aggregate Ethernet member links when the aggregated Ethernet bundle is part of an ECMP path. The aggregated Ethernet member links need to span the Virtual Chassis members. PR1255542
- The following cosmetic error is observed as the output: **mspmand[190]: msvcs_session_send: Plugin id 3 not present in the svc chain for session**. PR1258970
- When an interface comes online and both the OAM protocol and the MKA protocol tries to establish their respective sessions, OAM takes down the interface and MKA fails to establish connection (because the interface is down, it cannot send out MKA packets). PR1265352
- This very specific issue occurs when the Packet Forwarding Engine is oversubscribed with an unknown unicast flood with no MAC learning, which is not a common configuration. During unified ISSU, only the Packet Forwarding Engine gets wedged. However, this issue is not seen when the Packet Forwarding Engine is oversubscribed with Layer 3 traffic or with Layer 2 traffic with MAC learning. PR1265898
- On MPC2E-NG, MPC3E-NG, MPC5E, MPC6E, MPC7E, MPC8E, and MPC9E cards, a firewall performance feature **fast-lookup-filter** can be activated. Because of a transient parity error, the packet will be dropped within the PPE with the **sync xtxn error** message. This issue affects traffic, which might eventually affect the service. PR1266879
- GNFs in a node-slicing setup currently do not support Junos OS snapshot or recovery mechanisms. PR1268943
- Dynamic endpoint (DEP) does not support dh group19, encryption algorithm aes-256-cbc, and hash sha-384 in the list of default proposals. These must be configured explicitly in the configuration. PR1269160
- On a vMX platform, performance of the Intel X710 NIC is lower compared to the performance of Intel 82599 NIC. A 10-gigabit line rate can be achieved at 512 byte packet size for X710 NICs whereas the same can be achieved at 256 bytes for 82599 NICs. 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 from the previous VM host software instead booting from the alternate disk. PR1281554
- Due to a vendor code limitation, ungraceful removing of MX10003 MACsec TIC from the chassis might cause a crash or an unpredictable result. PR1284040

- This is in an internal change because syslog usage is deprecated; however, there might be customer impact due to syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful and relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. PR1284643
- This issue is observed on Junos OS Release 17.4R1-S3.3 image while testing the CUC-1422. **Error message: Jun 16 08:17:17 banaswadi rpd[51849]: Error creating dynamic logical interface from sub-unit 1051592: Device busy Jun 16 08:17:17 banaswadi rpd[51849]: Error creating dynamic logical interface from sub-unit 1051593: Device busy error message: rpd[51849]: Error creating dynamic logical interface from sub-unit 1051680: Device busy** PR1286042
- 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 line cards (MPC3E, MPC4E, MPC5E, MPC6E, MPC2E-NG, and MPC3E-NG) might report **DDR3 TEMP ALARM** chassisd's error log message. PR1293543
- A PCI device missing alarm might appear when the master and the backup Routing Engine run different versions of Junos OS. This PR adds the ability to verify if it is related to hardware or not before generating the alarm. PR1301191
- The lo0.0 interface should be used in default VRF for subscriber services. PR1303254
- Next-generation Routing Engine (NG-RE) with models RE-S-X6-64G, RE-S-2X00x6, and RE-PTX-X8-64G on the MX Series or PTX Series platforms might encounter a transient system freeze of the Linux-based host (VM Host) for about 20-35 seconds, causing protocol flaps, FPC restart, and mastership switch between the Routing Engines. Because of the incorrect handling of the disk I/O commands, a disk I/O timeout is reported and the system will recover by resetting the solid-state drives (SSD) channel. The system will continue to operate correctly after such an event. PR1312308
- The **show dynamic-tunnels database summary** command might not show an accurate tunnels summary during the time the anchor Packet Forwarding Engine line card is not in up state. As a workaround, use the following commands: **show dynamic-tunnels database** and **show dynamic-tunnels database terse**. PR1314763
- An alarm is raised if mixed AC PEMs are present. If the PEM is AC(HIGH) the first bit of pem_voltage is set, and if it is AC(LOW), second bit of pem_voltage is set. Therefore, if both the first and second bits are set, then MIXED AC is present. PR1315577
- The **chain-composite** configuration statement does not bring in a lot of gain because TCNH is based on ingress rewrite premise. Without this statement things work fine. PR1318984
- In JDM, (running on secondary server) jdmd daemon might generate a core file if GNF add-image is aborted by pressing Ctrl+c. PR1321803
- BGP signal 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
- In a streaming telemetry scenario, when performing **commit full**, the na-grpd daemon might restart, causing disconnection of streaming telemetry. PR1326366

- This is in an internal change as Syslog usage is deprecated, however, there might be a customer impact because of the syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful or relevant messages. PR1327266
- Per-protocol classification for host-bound traffic is not working. Because of the missing classification, protocol-specific (BGP/OSPF) statistics are not accounted correctly. PR1328631
- With regards to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infrastructure relies on the integrity of the TCP connections and the reactions to failure situations might not be handled gracefully. TCP connection timeout occurs because of jlock hog crossing the boundary value (5 seconds), causing bad consequences in the MX Series Virtual Chassis. Currently, no other easy solutions can reduce this jlock hog besides enabling marker infrastructure in the MX Series Virtual Chassis setup. Unfortunately, there is no immediate plan to enable marker because doing so causes several issues in the MX Series Virtual Chassis. PR1332765
- Under some race conditions with failover and multiple core interface flapping on an EVPN-VXLAN network, the rpd process might use high CPU, causing some issues in intercommunication with the l2ald process, and then causing the l2ald process to generate a core file and restart. PR1333823
- JFlow records are not sent out of the MS interface towards the collector unless we use SUBUNIT 0. PR1334682
- The output of the **show class-of-service fabric statistics** command now includes traffic that is dropped because of internal errors in the drop counts. PR1338647
- First packet pertaining to J-Flow Packet Forwarding Engine sensor in UDP mode is missing after line card reboots. PR1344755
- With GRES enabled in a subscriber environment, if subscribers are logging in or logging out very quickly, the sessions database in Session Database (SDB) of the backup Routing Engine might be leaked. If the problem is not detected for long enough, the backup Routing Engine might not be able to come back into synchronization with the master Routing Engine and will not be ready for GRES. PR1346300
- During 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 FPC restart. PR1348825
- On next generation Routing Engine (NG-RE), a failure of the hardware random number generator (HWRNG) will leave the system in a state where not enough entropy is available to operate. PR1349373
- In some cases, online insertion and removal (OIR) of a MIC on an FPC can lead to traffic that is destined to the FPC to be silently discarded. As a workaround, restart the FPC. PR1350103
- On all devices running Junos OS software, licenses might not take effect after successfully committing a license key configuration. PR1350302
- During stress conditions, error log messages regarding route add, change, or deletion might be incorrect. PR1350713
- When an ephemeral database instance is configured, if committing changes that are unrelated to IGMP/MLD (such as **set interfaces ge-0/0/1.0 description**), and the number of ephemeral commits reaches to the maximum size, the ephemeral database purge might happen. Then all the commits and

rollover might be purged. On this purge, the mgd gives all the applications a full commit view. And on this full commit view, IGMP/MLD deletes all the configurations and adds them back again. This might cause PIM to prune the groups on those interfaces and send join messages again. Finally, multicast traffic flapping and drop might be seen. PR1352499

- When performing unified ISSU to the Junos OS Release 18.2 and later releases with 1334612 fix, both the From and To builds should have the 1334612 fix. CRC errors are seen. PR1353911
- The issue occurs on aggregated Ethernet link deactivated or activated, that is the LAG interface is deleted from the system and created again. But then, the issue does not happen on deactivating or activating the link manually or by running this individual case in the script. There is no traffic loss. The traffic will continue to use the backup link. The aggregated Ethernet link up/down case is working as expected. Forwarding allocates a hardware selector for every primary link for local repair, which will be shared by multiple unilist next hops (a next hop with active and backup gateways using the primary and backup logical interfaces). The selector is getting stuck in rerouted state. There is no traffic loss but the traffic is flowing through the backup link even after the primary aggregated Ethernet link is created again. The problem seems to be with unilist->indirect->hold to unilist->indirect->unicast state transition during the deactivate or activate. As a workaround, enable the **vty** command to change the unilist hold behavior. PR1354786
- The craftd messages are generated on MX10003 and MX204 platforms. MX10003 platforms do not have 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 character in the fatal error message because the error no is not mapped to a string in the kernel code. *** **messages *** 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 might get fragmented. Because of this, at times the collector receives only the telemetry header part and not the payload. PR1364288
- After successfully delegating a locally configured LSP to a PCE, the router still displays 0 as the "Delegated" counter value under the output of CLI command **show path-computation-client status**. PR1369929
- The voltage high alarm might not be cleared when the voltage level comes back to normal for MIC on MPC5E. PR1370337
- Error messages are seen on an MPC card, but these errors do not impact any functionality on the card. **LOG: Err] PQ3_IIC(WR): bus transfer timeout on byte 1 LOG: Err] PQ3_IIC(WR): transfer not complete on byte 1 LOG: Err] PQ3_IIC(WR): I/O error (i2c_stat=0x21, i2c_ctl[0]=0xb0, bus_addr=0x76) LOG: Err] Failed to disable PCA9548(0x76)->channel(0-7) LOG: Err] zlpmb_set_channel: Failed to select channel 0 for MPC-PCIE1V0-LTC3880**. One of the root causes is that, the time to wait for the i2c transaction is not sufficient to finish the i2c transaction intermediately. Therefore, sometimes we see i2c transaction error. PR1374450
- Packet Forwarding Engine lookup loop happens when firewall-based redirection under **forwarding-options** is used to perform route-lookup in a non-default routing instance for destinations reachable MPLS over UDP tunnels. PR1378439

- In a subscriber scenario, if the **service-accounting-deferred** is configured on a dynamic profile, and there is multicast to a large number of destinations on the same physical port, FPC errors might be seen. PR1380566
- In case multiple LLDP sensors are getting exported together and part of their keys are overlapped, data for these sensors might not be exported. PR1382691
- Users can issue the command **set vmhost..** although **permissions system-control** is not configured on system class. PR1383706
- You can configure the purge timeout of programmable rpd clients to never. That is, 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 using 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
- On MX Series platforms enabled with a subscriber scenario, if a large scale of subscribers (for example, more than 1000 subscribers) set up connections simultaneously, the setup rate might be 30 percent lower than expected. PR1384722
- When traceoptions are enabled with a lot of trace flags or 'flag all', the rpd might crash because of the buffer overflow issue. This is a timing issue. PR1387050
- In low-end 32-bit systems, the rpd process has a lower level of available memory. It is desired to have a log message to alert the user when the average memory usage or transient memory usage exceeds thresholds. PR1387465
- During Zero Touch Provisioning (ZTP) process, the default route is being cleaned up by code. Because of this, If a static default route is configured in the initial configuration (configuration file downloaded from the file server for ZTP), the route might fail to work. This might lead to ZTP failure or a device access issue after ZTP. PR1387724
- On MX Series routers 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. PR1388120
- The bbe-smgd generates core files when MTU configuration is changed while subscribers are still logged in on the physical interface. MTU configuration change should only be done when there are no subscribers logged in on the physical interface. Catastrophic configuration changes should be done only in maintenance mode, when no subscribers are on the physical interface. PR1389611
- In Junos OS Releases after 17.3, ARP suppression and proxy ARP is enabled by default to reduce flooding of BUM traffic across the MPLS core files. But this is not certified in MPLS-VXLAN stitching scenario. When ARP suppression is enabled, in rare cases, when a local DC MAC entry timed out in VXLAN instance but no MPLS instance, the MPLS instance responds to the ARP request for that particular MAC address. This caused the VXLAN instance to install local MAC pointing to stitching point instead of VTEP/ESI next hop to local DC fabric. PR1390769

- 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 **set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress** command. PR1392071
- From Junos OS Release 16.1 onwards, the rpd might crash after executing the **show krt ack** command. PR1393959
- On MX2008 routers with MPC9E, in line-rate traffic with a redundant SFB2 scenario, if you take one redundant SFB2 offline, there might be tail or sometimes WRED drops in MPC9E, resulting in partial traffic loss. Under normal circumstances, the SFBs should be auto-failover if one of them fails, and there should be only a few packets dropped momentarily. PR1395591
- In a highly scaled EVPN-VXLAN environment, if there are many (1000+) simultaneous VM mobility events where the VMs move to reside behind a new leaf switch and the VM MAC addresses are also changed at the same time, in rare cases the ARP/ND table on the Layer 3 gateway devices might be left in stale state, pointing to the original leaf that hosted a VM rather than the new location. PR1395685
- MPC7, MPC8, or MPC9 cards have a local disk on which they keep a copy of the software image. The cards boot from the disk when an image is there, and boot from the chassis network (through BOOTP) when an image is not there. Presumably, new MPC7, MPC8, or MPC9 cards do not have an image on the disk and might require a network boot. On a single chassis, there is no problem. But on MX Series Virtual Chassis, the network boot does not work. PR1396268
- MPC card/afeb/tfeb with channelized OC MIC might crash and generate core files.PR1396538
- The rpd has facilities to attempt to trap certain classes of non-fatal bugs by continuing to run, but leaving a "soft" core file. Leaving a soft core files is intended to be non-disruptive to routing and forwarding. This PR implements a mechanism by which users might disable soft core files from being generated. PR1396935
- PPP lcp echo-request inline keepalives on an aggregated Ethernet interface are not sent after FPC is restarted. PR1397628
- The router advertises the ESMC quality level of the primary reference clock (PRC) even though the current clock status is holdover. PR1398129
- In a BGP-PIC case, if a route R1 resolves on top of a multipath-route R2, where R2 has primary and backup indirect next hops, it 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
- In a JET telemetry scenario, the telemetry log file is not rotated and keeps growing until the Routing Engine is out of disk space. This might cause an unexpected impact on the Routing Engine, and eventually lead to Routing Engine crash. The fix has now been provided to set the maximum allowable size to 50M and once the file reaches its maximum size, it will get rotated and compressed. PR1401817

- The authentication module for JET RPCs and telemetry fails in authenticating usernames or passwords of certain lengths. Hence the users will be unable to execute JET APIs or Junos Streaming Telemetry. PR1401854
- After upgrading to Junos OS Release 17.2 or later, the **chained-composite-next-hop ingress l3vpn extended-space** statement cannot be configured any longer on a logical system. PR1402390
- While initiating image installation on the base system of a setup with node slicing enabled, the session gets terminated unexpectedly. PR1402643
- Log messages similar to the following can be observed on MX10003 or MX10008 (a component might vary): **fpc0 mpcs_i2c_single_io: MPCS(0) ctlr 0 group 4 addr 0x4d prio 0 flags 0x10 failed status 0x1 fpc0 I2C Failed device: LTC3887-1:U15202 :EA1_VDD0V9R2, group 0x33 address 0x4d fpc0 pmbus_npcfpc_smb_write: npc_smb_write failed for group 0x33 addr 0x4d cmd 0x0 bytes 1 fpc0 ltc3880_pmbus_access_setup: Could not enable channel (channel 0) fpc0 pmbus_read_volt: summit-mx3ru-mpc - LTC3887-EA1-VDD0V9R2-CH0: pmbus read failed for cmd 0x8b fpc0 cmtfpc_pmbus_volt_get: Voltage read failed for rail LTC3887-EA1-VDD0V9R2-CH0, error : 1**. That is, a setting of a channel on i2c bus fails and hence a read operation fails after. It should not be an issue if subsequent readings are successful (no more messages are reported). This PR addresses the retry logic and tuning of the alerting mechanism. PR1405787
- On MX Series routers with MS-MPC card used, in race condition, if the MS-MPC is used on high availability (HA) scenario (the **set interfaces ms-x/x/x redundancy-options redundancy-peer/redundancy-local** statement and GRES is configured), the FPC might crash because of the bus error (segmentation fault). The reason is that when two CPUs simultaneously access the same session-extension memory in the session structure, one for writing, the other for reading. A reading CPU gets an incorrect value and uses that as the memory address. This causes the bus error (segmentation fault). PR1405917
- The rpd process might crash after a non-forwarding route (that is, a route to an indirect next hop association is a non-forwarding indirect next hop) which is received from multiple protocols is resolved again by using the non-forwarding path. PR1407408
- On MX Series routers using MPC7E, MPC8E, MPC9E, MX10K-LC2101, or MX10003, when in an inline-jflow application is used, a fatal error on Hybrid Memory Cube (HMC) performs a **disable-Packet Forwarding Engine** action. Because J-Flow records are hosted on the HMC memory partition, reading and writing to the HMC memory might trigger an FPC crash and high FPC CPU utilization. This causes slow convergence (adding/deleting routes or next hops) for other Packet Forwarding Engines on the same FPC carrier. PR1407506
- openconfig-network-instance:network-instances support for IS-IS is hidden unless supported. PR1408151
- In the scenario where **bgp multipath** is enabled, when forwarding chain is unilist_1->indirect-next-hop->unilist_2, any change in unilist_2 active member list will be absorbed by indirect-next-hop in the chain and the change will not be back propagated to top-level unilist_1. If a link flaps it will cause indirect-next-hop pointing to unilist_2 stuck with weight 65535 and further causing traffic getting silently dropped and discarded. PR1409632
- The configuration database can remain locked after the SSH session is halted. PR1410322

- In highly loaded subscriber management setups, some SNMP queries might experience a response delay from the MX Series router due to higher priority daemons utilizing CPU resources. PR1411062
- If a 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 MPC8 line card, enabling both bandwidth configuration statement along with **flex-flow-sizing** statement might result in J-Fflow service getting disabled because unable to allocate the memory requested by **flex-flow-sizing** statement. PR1413513
- 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
- On MX Series platforms, if non-default MTU (for example, 4400) is configured on PS physical interface, when performing a GRES or dcd restarts, the dcd triggers catastrophic events below the IFF (interface family). This might cause deletion and addition of IFAs (interface address) and it causes protocol sessions (such as BGP session) on this PS interface to flap. PR1415207
- PCE initiated LSPs get deleted from PCC if the PCEP session goes down and gets reestablished within the **delegation-cleanup-timeout** period. PR1415224
- With NETCONF, the xmlns attribute is printed twice for RPC get-arp-table-information to the router. PR1417269
- In a DHCP subscriber scenario, some subscribers might be offline when doing GRES or daemon restart. The reason is that when restoring the subscribers back from the subscriber database (SDB), the profile database (PDB) call (an internal call that is done to determine the interface type) fails. Sometimes PDB calls are unreliable and could return an error if the database is not ready. This is also the root cause. PR1417574
- PTP/hybrid is not supported with the hyper mode. Delete or deactivate hyper-mode configuration and reboot the router to use PTP #delete forwarding-options hyper-mode or add the below configurations on the platforms where supported **set forwarding-options no-hyper-mode**. PR1420809
- PF Core Voltage is not set as per the required e-fuse value and remains to default value (0.9V). PR1420864
- On MX Series routers with 1xCOC12 or 4XCOC3 used, if channelized interfaces are configured, FPC CPU hog might be seen. PR1420983
- The XML formatted output of command **show security group-vpn member ipsec statistics** is not hierarchically structured which does not allow to easily associate <esp-statistics> and <ah-statistics> elements with the respective <usp-ipsec-service-set-statistics> elements.PR1422496
- For MX204 routers, the number of PICs per FPC is incorrectly used as 8, which causes a MAC allocation failure on the physical interfaces. PR1422679
- **error: mustd trace init failed** during configuration commit PR1423229

- Configure commit error is seen when **dhcp traceoption** is disabled. **error: DHCP service may not be de-configured while clients are present. Please clear bindings**. PR1423500
- On Junos OS routers and switches with link aggregation control protocol (LACP) enabled, deactivating a remote aggregated Ethernet member link will make the local member link move to LACP detached state. The detached link will be invalidated from the Packet Forwarding Engine aggregated Ethernet forwarding table as expected. However, if the device is rebooted with this state, all the member links will be enabled in Packet Forwarding Engine aggregated Ethernet forwarding table irrespective of LACP states and result in traffic drop. PR1423707
- 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 vMX platforms, the link flapping for the ixgbe interface might trigger the physical function (PF) to reset for IXGBE, but the virtual function (VF) reset will not be done. The issue results in traffic drop for the interface. PR1424626
- The issue is limited to database and is related to MAC-MOVE scenario. When **dhcp-security** is configured, if multiple IPv4 and IPv6 client's MAC-MOVE happens, the jdhcpd might consume 100 percent CPU and jdhcpd crashes afterwards. PR1425206
- The rpd crashes on routers running a 64-bit VM host image if protocol authentication is used along with the master password. Any routing protocol that supports authentication can trigger this; some examples include BGP, IS-IS, and OSPF. PR1425231
- Whenever the **show snmp mib walk jnxMibs** command is issued, the below logs are observed in the chassisd: **Mar 14 15:59:33 fru_is_present: out of range slot 0 for Mar 14 15:59:33 fpm_get_sys_led: FPM display module missing Mar 14 15:59:33 snmp_get_pem_led_state 936: pem state = 5, ret_val = 2 Mar 14 15:59:33 snmp_get_pem_led_state 936: pem state = 5, ret_val = 2**. The above logs are triggered by SNMP polling. These logs are superficial in nature and have no impact on production. Please refer to: https://kb.juniper.net/InfoCenter/index?page=content=KB24394. PR1425411
- On a port-down event for an L2BSA subscriber, the calculated rate might be set as 0 and this calculated rate as 0 might be sent to RADIUS over accounting stop message. PR1425512
- Because of misconfiguration (overlap in the lo0 address and subscriber address), causing an incorrect reference count for the local address and the local route gets stuck in bbe-smgd. The fix checks for the route type to ensure that the reference count is not increased in this case. If the subscriber address matches with the local address, the subscriber goes offline. PR1428428

- Finisar optic is reporting oncorrect threshold information for Rx power. Therefore Junos OS will not set the alarm when the signal goes below the manufacturer specification. PR1430842
- On MX Series platform with FPCs, if **sa-multicast** is configured, all the traffic get dropped. PR1433306
- Core might be generated on the backup Routing Engine when virtual router and forwarding instances configuration with NSR are enabled. On MX Series routers the following log messages are seen: **MX@ 0x0000000001662f76 in rpd_name_cmp (a1=0x489046c, a2=0xffffffff04b73cb8)** at **../../../../../../../src/junos/usr.sbin/rpd/lib/common/rpd_vrf_cmp.c:45**. PR1433883
- Interoperable issues exists with INNOLIGHT QSPFs (Part number 740-054050) that cause link flap. PR1436275
- There are some corner cases or race conditions where when a delinked marker is in the IFSTATE chain, the subsequent ifstate might get missed. This might keep on lingering in the ifstate chain and the particular client (mib2d in this case) will never be able to read it. This causes a high level of CPU utilization. As a workaround, restart the (mib2d) daemon from the CLI. PR1437762
- On all Junos OS platforms, if the hash-key is enabled, packets might be dropped because the chassisd crashes and even packets on other FPCs, the hash-key is disabled. PR1437855
- SNMP trap comes twice for FRU removal in MX10000 one trap with FRU name as FPC: JNP10K-LC2101 and second with FRU name as FPC @ 1/*/*. PR1441857

Infrastructure

- A file system corruption might create a kernel core file. The Routing Engine reboots with the message **ffs_blkfree: freeing free block**. PR1028972
- The following messages are seen during FTP: **ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock (No such file or directory)**. PR1315605
- All devices 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. It should be later than 1.1 for this warning to be resolved. PR1345166
- 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
- With **master-password** configured, a software upgrade might fail with a message **Hardware Database regeneration succeeded Validating against /config/juniper.conf.gz /config/juniper.conf:833:(199) secret: invalid encoded string at '8aes256-gcm$hmac-sha2-256$100$Ue0Pb8XNWS0$**. PR1404431
- With newer Junos OS images increasing in size, the current partition scheme for the CF is not allowing enough room for restoring a full recovery snapshot. This is mitigated by temporarily disabling the swap file during recovery process. PR1420356

Interfaces and Chassis

- Junos OS now checks logical interface information under the aggregated Ethernet interface and prints the information only if it is part of it. PR1114110
- In a VPLS multihoming scenario, the CFM packets are forwarded over the standby PE device link, resulting in duplicate packets or a loop between the active and standby links. PR1253542
- Out-of-sequence packets are seen with the LSQ interface. PR1258258
- In Junos OS BNG solutions, after a commit event, when the configuration contains **duplicate vlan-id** configured on aggregated Ethernet and demux interfaces, the MX Series router might go into database prompt mode and the kernel generates core files. PR1274038
- Upgrading to 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 presence of an old version of **/var/db/cfm.db** file. PR1281073
- LAG member links running LACP in slow mode might get disassociated from the LAG bundle with a combination of restart interface-control and FPC offline/online trigger. The issue is seen with scale configuration on DUT. The scale details are 2800 CFM sessions, 2800 BFD sessions, 2043 BGP peers, and 3400 VRF instances. PR1298985
- The CFM session does not come up if configured on a logical interface with a VLAN ID matching that of the configured native VLAN ID under the physical interface. PR1325190
- If 64,000 bridge domains are configured, with each BD having two logical interfaces and 1 IRB interface, heap memory exhaustion occurs as this requires more than the supported memory on the FPC. The possible workaround for this is to configure interfaces in the trunk mode that allow all 4000 VLANs, reducing the need to configure the logical interfaces for each BD. The trunk ports are configured in the default instance or for each routing instance. PR1348363
- In MX Series Virtual Chassis, flooding of the error message **CHASSISD_CONFIG_ACCESS_ERROR: pic_parse_ifname: Check fpc rnage failed** can be seen with LACP-enabled aggregated Ethernet interfaces on MPC7, MPC8, or MPC9 cards. Errors will only have impact for DWDM PICs, which does not have an effect on the MPC7, MPC8, or MPC9 cards. Hence, this syslog message can be safely suppressed. PR1349277
- The **ppman_cfm_start_inline_adj: Failed to add Inline adj for CFM, pkt-len=0** error message will be observed in some cases. But there is no functional impact. Sessions/adjacency might get programmed inline subsequently. PR1358236
- There might be memory leaks on transportd when bulk SNMP polling done are on large-scale logical interfaces and a large number of traps are created because of interface flapping. The memory leak might cause the transportd to consume high CPU for a prolonged period. PR1398967
- Static demux0 logical interfaces do not come up after configuration change if the underlying interface is et (100-gigabit Ethernet). After a configuration change, the et interface gets flushed in order to reparse the configuration. During this, 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 already been addressed. This is day-one issue. As a workaround, restart the DCD (or reboot the entire Routing Engine), to clear the problem or else use **commit full** instead of **commit** while committing the new configuration. PR1401026

- On MX Series routers, EX-SFP-1FE-LX SFP does not initialize with MIC-3D-20GE-SFP-E(EH). PR1405271
- ICCP does not come up when mc-lag PE is rebooted since static ARP is deleted and never re-installed back. So it is not recommended to configure ICCP over IRB which is associated with mc-lag bridge domain. Customer upgrading from old release to new release (PR 1075917 support) might come across issue like static ARP not is not reinstalled for remote mc-lag IRB IP when existing static ARP entry is removed. PR1409508
- When an unnumbered interface is binding to an interface which has more than one IP address and one of the IPs is deleted, the family inet of the unnumbered interface might be getting 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, thereby avoiding the deletion of the family inet of the unnumbered interface. PR1412534
- If the aggregated Ethernet interface has VRRP configuration, in following use cases, member logical interface will not be created after member physical interfaces comes up and aggregated Ethernet will be in down state. 1. FPC restart (**request chassis fpc restart slot <>**) 2. chassis-control restart (**restart chassis-control**) 3. reboot both Routing Engines (**request system reboot both-Routing-Engines**). So, before performing above operations, it is advisable to remove VRRP configuration from aggregated Ethernet interface. PR1429045
- Because of a software issue, aggregated Ethernet interface outgoing traffic might get dropped on all ingress Packet Forwarding Engines by hitting sw.nh.discard_sampling_trap. PR1441772

Layer 2 Ethernet Services

- In an MC-LAG with force-up scenario, LACP PDU loop might be seen when both MC-LAG nodes and the access device use the same admin key. PR1379022
- On EVPN setups, incorrect destination MAC addresses starting with 45 might show up when using the **show arp hostname** command. This is a cosmetic issue with no impact. PR1392575
- On MX5, MX10, MX40, MX80, and MX104 routers with DHCP server configuration for DHCP subscribers, the jdhcpd memory leak might happen and the memory increases by 15 MB, which depends on the number of subscribers when testing the DHCP subscribers login or logout. PR1432162

Layer 2 Features

- For a router equipped with the following line cards: T4000-FPC5-3D, MX-MPC3E-3D, MPC5E-40G10G, MPC5EQ-40G10G, or MPC6E MX2K-MPC6E line cards, if the router is working as a VPLS PE device,

because of MAC aging every 5 minutes, the VPLS unicast traffic is flooded as unknown unicast every 5 minutes. PR1148971

- In an LDP-VPLS setup where user-defined mesh groups are configured in a VPLS instance and the LDP-VPLS must also have at least one directly connected CE interface configured under the instance, and 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. And if **connectivity-type permanent** is configured, this issue will not be observed as the instance will remain UP state. PR1415522

MPLS

- When using the **mpls traffic-engineering bgp-igp-both-ribs** statement with LDP and RSVP both enabled, CSPF for interdomain RSVP LSPs cannot find the exit area border router (ABR) when there are two or more such area border routers (ABRs). This causes interdomain RSVP LSPs to break. RSVP LSPs within the same area are not affected. As a workaround, you can either run only RSVP on OSPF ABR or IS-IS L1/L2 routers and switch RSVP off on other OSPF area 0/IS-IS L2 routers, or avoid LDP completely and use only RSVP. PR1048560
- The issue occurs when GRES is done between the master and backup Routing Engines of different memory capabilities. For example, one Routing Engine has only enough memory to run routing protocol process (rpd) in 32-bit mode while the other is capable of 64-bit mode. The situation might be caused by using Junos OS Release 13.3 or later with the configuration statement **auto-64-bit** configured, or, by using Junos OS Release 15.1 or later even without the configuration statement. Under these conditions, the rpd might crash on the new master Routing Engine. As a workaround, this issue can be avoided by using the **set system processes routing force-32-bit** CLI command. PR1141728
- 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 (say FPC containing both primary and bypass or, both primary and bypass FPCs go down simultaneously) such that primary link goes down even before the PLR sends out any path message after bypass 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 Explicit Route Object (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
- An SR-TE path with "0" explicit NULL as inner most label, 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
- Packets loss might be observed when auto-bandwidth is enabled for circuit cross connect (CCC) connections and label switched path (LSP) **no-self-ping** with **no-install-to-address** is configured. PR1328129
- Executing a **restart chassisd** in a router with a scaled configuration might result in generating rpd core file. PR1352227
- If family MPLS is disabled or deleted from the interface and enabled or added back to the same interface or the maximum number of labels are changed in quick succession, without a delay, the routes might be dead in the forwarding table with a label using this interface as out interface. As a workaround, it is recommended to add a slight delay (5 seconds is recommended) for this issue. PR1355878
- Traceroute MPLS from Juniper Networks routers 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 **traceroute mpls bgp** command, in very rare cases, it is possible that 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 might cause traffic impact till rpd process comes back up. PR1399484
- With NSR enabled, when master rpd process is restarted, occasionally, out-of-order add and delete messages can arrive on the backup Routing Engine, causing label assignment collisions leading the backup rpd to crash. PR1401813
- When make-before-break (MBB) new instance signaling experiences error and before retry is finished, other triggers such as auto bandwidth adjustment timer expiration have to be blocked until MBB finishes. After MBB switches instances, 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 **optimize-adaptive-teardown** timer into account. This causes the old instance to be torn down before the whole system finishes changing routes using the new instance, that leads to traffic loss. PR1402382
- On Junos OS devices, 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 percent for hours. The issue might affect traffic. PR1405033
- In a BGP-labeled unicast scenario with **egress-protection** enabled, the rpd process might crash while committing some configuration changes even if the changes are not related to **egress-protection** itself. PR1412829
- The LDP transit egress route for a BGP route has an indirect next-hop. In NSR and GRES scenario, after Routing Engine switchover, the LDP might fail to receive the route flash for a BGP route from inet.0 and might not update the inet.3 route for the BGP route. As a result, the next hop for LDP transit egress

route becomes unusable and the LDP transit egress route might get deleted. As a result, BGP sessions go down and traffic drops. PR1420103

- LDP route metric might not match IGP route metric even with **ldp track-igp-metric** configured. PR1422645
- When current statistics collected are lower than previous values, it is not correct to consider counters overflow and use current statistics to calculate rate. Because of the incorrect assumption, the rate can be skewed. There can be various reasons, including route changes and next hop changes to cause temporary statistics decrease than last received statistics, when this happens, the rate calculation should be avoided and update statistics only to prepare for next statistics collection and rate calculation. PR1427414
- The rpd process core files might get generated on backup router during label allocation after performing GRES operation. PR1427539
- Dynamically configured RSVP LSPs for LDP link protection might not come up after disabling or enabling protocol MPLS. PR1432138
- In LDP to BGP-LU stitching scenario, when user performs a MPLS ping, if the BGP routes for that LSP is unavailable, the rpd process will crash. PR1436373

Network Management and Monitoring

- A vulnerability in Junos OS SNMP MIB-II subagent daemon (mib2d) might allow a remote network based attacker to cause the mib2d process to crash resulting in a denial of service condition (DoS) for the SNMP subsystem. PR1241134
- 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 **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 corresponding freeing for this structure is not done, which causes the memory leak. PR1392616

Platform and Infrastructure

- In configurations with IRB interfaces, during times of interface deletion (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. PR1054798
- The logs such as **cassis_alloc_index_pool_create: SVC NH 0x00b00000[0] poolsize 0x000fffc0 is not a multiple of blk_sz 0x00001000**. The logs are cosmetic and has no service impact. PR1301924
- 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

- On all devices running Junos OS, execution of Python scripts through enhanced automation does not work on veriexec images. PR1334425
- This is a minor enhancement to add a UI to copy files from Junos OS VM to host Linux. PR1341550
- In a filter list (input-list or 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 PE device for a CE-facing network results in an ICMP TTL expired reply with a source address of only one of the many CE-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. PR1358376
- BBE CST telenorSweden MPC7 is stuck in ready state and CPU usage is 100 percent after unified ISSU from Junos OS Release 17.3R2-S4 to 17.3R2-S4 forever. PR1368854
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps happens even with **enhanced convergence** configured. There is no guarantee that subsecond convergence will be achieved. PR1371493
- In a Layer 3 VPN network with large-scale prefixes, if the peer PE device is another vendor's router configured with **per-prefix label**, all FPC cards might restart after the Layer 3 VPN routes churn multiple times. PR1398502
- In a subscriber management environment with scaled subscribers login such as 200k PPPoE subscribers, FPC crash might be observed. PR1409879
- On MX Series platforms with VPLS scenario, when the **interface-mac-limit packet-action-drop** statement is configured, in the case of MAC moves, the new MAC might not be learned sometimes because of a race condition of an unusual update of "MAC learn limit" under the Packet Forwarding Engine (the hardware "MAC learn limit exceeded" counter displays unexpected behavior and increases to a very huge and negative number). This can result in packet drops. PR1410162
- On MX Series with MPC, if an aggregated Ethernet interface is with the filter of **shared-bandwidth-policer** and the **shared-bandwidth-policer** statement is deactivated, after activating the **shared-bandwidth-policer** statement, the policer bandwidth might be calculated as 0 and all traffic might be dropped for the aggregated Ethernet interface. PR1427936
- In some scenarios with EVPN, customer might experience the following messages, followed by PPE traps on **/var/tmp** showing **Async XTXN Error PPE/Context 33/27 @ PC 0x076f: dmac_miss_read_bd mx1 fpc0 XL[0:0]_PPE 33.xss[0] ADDR Error. mx1 fpc0 XL[0:0]_PPE 33 Errors async xtxn error**. These are isolated entries and there is no service impact. PR1428456
- On MX Series platforms enabled with GRES and nonstop Routing (NSR) on both end of BGP peer, in scaled BGP session setup, BGP peers might flap after the execution of Routing Engine mastership switchover on both the boxes simultaneously. PR1437257
- The CLI displays incorrect next hop MAC address in **show route forwarding table** command. PR1437302

Routing Policy and Firewall Filters

- The rpd process might crash during the policy configuration changes. PR1357802
- If a **policy-option** with only conditions from route-distinguisher and then next-hop a.b.c.d is applied to BGP, the next hop for routes in the inet.0 might be set to this next-hop a.b.c.d, even though these routes do not carry any route distinguisher value (l3vpn.inet.0 is unaffected). PR1433615

Routing Protocols

- When only the default routing instance is present, the **show bgp summary** command does not show the BGP ESTABLISH state. If the BGP state is not an ESTABLISH state, then it shows the states as design (that is active, idle, connect). If there is a routing instance configured (apart from master routing-instance inet.0), the BGP ESTABLISH state is showed properly. The issue happens for IPv4 BGP sessions only, on IPv6, all the BGP states are shown as default. PR600308
- In rare cases, rpd process might generate a core file with **error rt_notbest_sanity: Path selection failure on**. But, there might be no impact to traffic or routing protocols. PR946415
- When interoperating with other vendors in a draft-rosen multicast VPN, by default the Junos OS attaches a route target to the multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities will be excluded from propagating if the BGP route target filtering is enabled on Junos OS device. PR993870
- The static/static access routes pointing to an unnumbered interface are getting added in the routing table even if the interface is down. In this case, if GRES is disabled, this type of route will never be added in the routing table after Routing Engine switchover. PR1064331
- The syslog message **JTASK_SCHED_SLIP** for rpd process might be seen on restarting routing or disabling ospf protocol with scaled BGP routes in the MX104 router. PR1203979
- Certain BGP traceoption flags (for example, "open", "update", and "keepalive") might result in (trace) logging of debugging messages that do not fall within the specified traceoption category. This results in some unwanted BGP debug messages being logged to the BGP traceoption file. PR1252294
- LDP OSPF are 'in sync' state and the reason observed for this is **IGP interface down** with **LDP-synchronization** statement enabled for OSPF: **user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.0 1 Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050 Adj count: 1 Hello: 10, Dead: 40, ReXmit: 2, Not Stub Auth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTC Protection type: None Topology default (ID 0) -> Cost: 1050 LDP sync state: in sync, for: 00:04:03, reason: IGP interface down config holdtime: infinity**. As per the current analysis, **IGP interface down** is observed as the reason because although LDP notified OSPF that LDP synchronization is achieved, OSPF is not able to take note of the LDP synchronization notification, because the OSPF neighbor is not up yet. PR1256434

- This is in an internal change as syslog usage is deprecated. However, there might be a customer impact because of syslog usage in automation. Applications have migrated to tracing for engineering debug messages or ERRMSG for customer useful or relevant messages. The customer is advised to migrate to new ERRMSG definitions as appropriate. PR1284621
- When the **clear validation database** command is issued back-to- back multiple times, it results in partial validation database. This eventually recovered after up to 30 minutes (half of the record lifetime) when periodical full updates are done. PR1326256
- In a large-scale OSPF network (for example, if there are more than 500 devices in an area), OSPF remote loop free alternate (rLFA) default PQ node selection algorithm does not provide proper protection paths. PR1335570
- There are scenario where application allocates and caches next hop templates. This causes next hop template cache to grow continuously. But when application clears their local cache, then memory is freed to next hop template cache. But the next hop template cache does not have 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 cache and not consume additional memory. PR1346984
- Traffic loss can be seen with Layer 3 VPN egress protection after GRES. PR1347980
- When the loopback interface is configured in a logical system and the Routing Engine based micro BFD is configured to use the loopback address as source address, BFD packets go out with the source address belonging to the outgoing interface rather than the loopback address. Due to this issue, the micro-BFD session might not be able to come up. PR1370463
- Committing the configuration might fail on the backup Routing Engine in certain conditions. PR1372847
- When the OSPF segment routing one-hop neighbors are connected through the unnumbered interface, an invalid label operation might happen. PR1386133
- In a BGP scenario with multipath enabled, when applying the import/export policy of IPv6 routes with an IPv4 next hop to a BGP neighbor, the rpd process might crash continuously. PR1390428
- If an import policy is applied to a BGP neighbor and the policy has indirect IPv4 next hop for IPv4 and IPv6 routes (IPv6 routes resolved over IPv4), when BGP unresolved route is withdrawn, the rpd process might crash. PR1391568
- The **as-path-group** configuration is limited in scale. With 10000 lines, scheduler slips are seen, impacting other work the rpd process does such as protocol keepalives. To avoid the scheduler slips (CPU exhaustion), change how the as-path-group is structured. The issue occurs because of the two factors - the number of as-path statements under the as-path-group and the wildcards in each of these. PR1396344
- 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 suboptimal handling of MoFRR forwarding states on the Packet Forwarding Engine level. PR1399457

- 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 the **qualified-next-hop** points to the gateway family data links (GF_DLI) address, which the PIM is unable to process, resulting in the crash. PR1408443
- When a BFD session timeout, OSPF sends an LS update. If the BFD peer router receives the LS update through alternative OSPF path BFD timer might be refreshed. This causes a delay in detection the neighbor loss. This behavior is seen in both IPv4 OSPF BFD sessions and IPv6 OSPF BFD sessions in case of using centralized BFD. PR1410021
- In a BGP with the indirect next-hop scenario, if uRPF 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 percent. PR1414021
- In MVPN scenario, the rpd process might crash while removing multicast routes that do not have an associated (S,G) state or activating the **accept-remote-source** statement on PIM upstream interface. PR1426921
- In a BGP graceful restart scenario, including helper mode which is enabled by default, rpd process might generate core files because of the improper handling of BGP graceful restart stale routes when the BGP neighbor is deleted. The rpd crashes and there is service or traffic impact. PR1427987
- In a BGP Labeled Unicast (BGP-LU) scenario, if the device works as penultimate hop and receives BGP-LU routes with indirect next-hop from an egress router, after the operational next-hop interface corresponding to those labeled routes flaps, a "dead" next-hop type (discard action is performed for this type) may be set for the related clone routes (s=0) and still there even the next-hop interface is operational again. The issue present again only in 16.1R7-S3 after PR1333570, and the fix is complete in 16.1R7-S5. PR1432100

Services Applications

- Hide HA information when the service set does not have ha configured. PR1383898

Subscriber Access Management

- Sometimes, when PPPoE subscribers log in and log out from Junos OS Release 16.1, 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 authd daemon for subscriber authentication is attempting to read private data for an underlying interface which no longer exists (-7 = SDB_DATA_NOT_FOUND).** These messages have no impact and can be safely ignored, where authd daemon is asking sdb for record that no longer exists. PR1236211
- On any JunoS platform running the authd process, the dot1x user might fail to authenticate with the Radius server if the connection to Radius server bounced. PR1338993

- The authd reuses addresses too quickly before jdhcpd completely cleans up the old subscriber which causes flooding of error logs 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
- In a chain of linked address pools, if the last pool is sent and linked pool aggregation is configured, the device does not return to the linked pool once the last pool is consumed. PR1426244

User Interface and Configuration

- The mustd generates core file ppool_bkt (phdr=0xde918024, pfile=0xde933004, no_pages=1) at **../../../../../src/ui/lib/memory/page_pool.c**. PR1309074
- The test configuration **/config/rescue.conf.gz** fails commit check for a dynamic profile when the subscriber is active. PR1376689
- The **show chassis hardware satellite** command is not available on Junos OS Release 17.3. **root@mx104> show chassis satellite detail Satellite Alias: fusion FPC Slot: 101 Operational State: Online <...> Below, you can see no "show chassis hardware satellite" option: root@MX104> show chassis hardware ? Possible completions: <[Enter]>** Execute this command clei-models Display CLEI barcode and model number for orderable FRUs detail Include RAM and disk information in output extensive Display ID EEPROM information models Display serial number and model number for orderable FRUs | Pipe through a command PR1388252

VPNs

- In an MVPN environment with SPT-only option, if the source or receiver is connected directly to c-rp PE and the MVPN data packets arrive at the c-rp PE before its transition to SPT, the MVPN data packets might be dropped. PR1223434
- The multicast VPN MIB is not properly compiled into the Juniper Networks MIB package bundle. This PR causes **mib-jnx-mvpn.txt** to be included as part of the Juniper Networks MIB set. PR1394946
- When the end-interface or backup-interface/protect-interface in the end-interface is used as an interface for the **ping mpls l2circuit interface** command, the rpd process might crash and generate core files. PR1425828

SEE ALSO

New and Changed Features | **112**

Changes in Behavior and Syntax | **134**

Known Behavior | **145**

Resolved Issues | **176**

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

Application Layer Gateways (ALGs)

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

Authentication and Access Control

- 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)

- The cosd process might crash during committing configuration change through netconfiguration. PR1403147
- Traffic drop occurs when deleting MPLS family or disabling interface which has non-default EXP rewrite-rules. PR1408817

Flow-Based and Packet-based Processing

- When Power Mode IPsec feature is enabled, and fragmented traffic is received by the SRX on an IPsec tunnel, the tunnel is moved from Power Mode IPsec to regular flow IPSec mode. PR1397742

EVPN

- L2ALD restarts when changing "protocols" related configuration. PR1357911
- The EVPN implementation does not follow RFC-7432. PR1367766
- The rpd process might crash when deactivating the autonomous system (AS) in an EVPN scenario. PR1381940
- The RA packets might be sent out without using the configured virtual gateway address. PR1384574
- [EVPN/VXLAN] VTEP tunnel does not get deleted when EVPN peer goes down. PR1390965
- A few minutes of traffic loss might be observed during recovery from link failure. PR1396597
- The BUM traffic might not be flooded in EVPN-MPLS scenario. PR1397325
- IPv6 link-local address for virtual gateway address is marked as duplicate in EVPN. PR1397925
- When committing a configuration for a VLAN adding 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 two MPLS labels and withdrawal has one label. PR1399726
- The rpd process generates a core file upon Routing Engine switchover with scaled EVPN configuration. PR1401669
- The rpd crashes because of the memory corruption in EVPN. PR1404351
- EVPN database and bridge mac-table are out of synchronization because of the interface 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
- EVPN routes might show **Route Label: 0** in addition to the real label. PR1405695
- The rpd might crash after NSR switchover in EVPN scenario. PR1408749
- Local L2ALD proxy MAC and IP advertisements accidentally delete MAC and IP EVPN database state from remotely learned type 2 routes. PR1415277
- The rpd crashes on backup Routing Engine after enabling nonstop-routing with EVPN. PR1425687
- The CE interface IP address is missed in **mac-ip-table** of the EVPN database. PR1428581
- Stale MAC addresses are present in the **bridge mac-table** in EVPN or MPLS scenario. PR1432702

Forwarding and Sampling

- Firewall **from packet-length** match with more than two ranges will fail on PE/TL for QFX Series platforms without warning. PR1221777
- The kernel crash might be observed when there is a firewall filter modification PR1365265
- In EVPN A-A scenario with MX or EX acting as 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
- LTS subscriber statistics are reported to RADIUS. PR1383354
- The LSI binding for the IPv6 neighbor is missing. PR1388454
- When **flex-match-range** functionality is configured using bit-length 128 for IPv6, the errors with configuration settings are seen. PR1389103
- On Junos OS the firewall filter terms named "internal-1" and "internal-2" are ignored. PR1394922
- 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 is seen in the syslog. PR1232178
- We advise migrating from syslog API to Errmsg API: **/src/junos/usr.sbin/mspsmd**. PR1284654
- Large-scale users logging in and loging 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 a master Routing Engine switchover in a VRRP scenario. PR1353583
- On MX Series routers, network slicing GNF is allowed to install incompatible images without warnings. PR1353773
- Packets might be dropped when they go through the MX104 built-in interface. PR1356657
- FPC core files might be generated after GRES switchover. PR1361015
- On the MX10003, the alarm LED reflects the stale entry on the backup Routing Engine, post GRES switchover. PR1361728
- On Junos OS, the **set system ports console insecure** command allows root password recovery on OAM volumes. PR1368998
- The MPC5E, MPC2E-NG, and MPC3E-NG might crash and restart during unified ISSU. PR1369635
- The filter service might fail to get installed for the subscriber in a scaled BBE scenario. PR1374248
- MS-MPC might have performance degradation under scaled fragmented packets. PR1376060
- The **Power Supply failed** trap might not be generated on MX Series routers. PR1376612

- Traffic loss might occur during unified ISSU on MX Series routers with MPC7E, MPC8E, and MPC9E. PR1377782
- MQTT keepalive timeout messages are seen in case of slow JTI collectors. PR1378587
- The routes learned over an interface might be marked as "dead" next hop after the prefix length of an IPv6 address is changed from 128 to any other prefix length. PR1380600
- Traffic gets silently drops and gets discarded without a notification when an FPC is taken down in an MC-LAG scenario. PR1381446
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent the major alarm. PR1384435
- The device with more than five IP addresses configured in the DHCP server group goes into amnesiac mode after reboot. PR1385902
- The rpd and KRT queue might get stuck in a VRF scenario. PR1386475
- On MX2000 routers, backup CB's chassis environment status keeps 'Testing' after backup CB becomes online by removal/insert operation. PR1387130
- Some SFBs might go down when one of the PSMs in the chassis generates a bad output voltage which is out-of-range. PR1387737
- FPC core file is seen at **sensor_export_time_exceed_limit agent_health_monitor_data_reap** when **Jinsight** is configured. PR1388112
- IPsec IKE keys are not cleared when delete or clear notification is received from peer on GRES-enabled DUT. PR1388290
- The jnxFruState might show incorrect PIC state after replacing an MPC with another MPC having less PICs. PR1390016
- Traffic destined to VRRP VIP gets dropped as filter is not updated to related logical interface. PR1390367
- The BNG might not respond with PADO and create any demux interface when PPPoE PADI packet is received. PR1390989
- The Packet Forwarding Engine might not respond with ICMP time exceeded error when packets arrive from subscribers. PR1391932
- Third-generation FPC reboot loops because of internal intf issues. PR1393643
- FPC might reboot on vMX in a subscriber scenario. PR1393660
- Junos OS enhancement configuration statement added to modify mcontrol watchdog timeout. PR1393716
- The FPC cards might not come up while performing unified ISSU on MX10003. PR1393940
- The gRPC hardcoded credentials might allow unauthorized access to systems with Junos Network Agent installed. PR1394927
- On the MX960 router, a minor alarm “Bottom Fan Tray Pred Fail” might incorrectly be raised while fans are running at high or full speed. PR1395539

- Adding IRB to a bridge domain with a PS interface causes kernel crash. PR1396772
- On MX1008, VMHost **VMHost RE 0 Secure BIOS Version Mismatch** and **VMHost RE 1 Secure Boot Disabled** alarms are seen. PR1397030
- On PFT MX10008 routers, the inline-services enabling the **Flex-Flow-Sizing** takes more than 12 minutes to move to steady state. PR1397767
- The **show system firmware** command might provide unexpected output on some MX Series platform such as MX104. PR1398022
- The **show system errors active** command is not showing the error for MPC3E NG HQoS. PR1398084
- IPsec tunnel cannot be established because that the tunnel SA and rule are not installed in the PIC. PR1398849
- Incorrect timestamp is displayed in the Junos Telemetry Interface collector log file. PR1399829
- The mgd-api might crash because of a memory leak. PR1400597
- Only one Packet Forwarding Engine could be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- The Framed-Route beyond the first might not be installed in a DHCP subscriber management environment. PR1401148
- The authd process might crash when issuing the **show network-access requests pending** command during the authd restarting. PR1401249
- The **show | compare** command output on global group changes lose the different context after a rollback or **load update** is performed. PR1401505
- The subscriber route installation fails because some interfaces states are not properly installed. PR1401506
- The TCP connection between ppmd and ppman might be dropped because of the kernel issue. PR1401507
- Fabric packet forwarding performance might be degraded on MX2000 routers installed with SFB2 and MPC9E. PR1401599
- FPC core files are generated because of a corner case scenario (race condition between RPF, IP flow). PR1401808
- Traffic loss is seen in IGMP subscribers after GRES. PR1402342
- The MPC might crash because of the CPU hogging by dfw thread. PR1402345
- Some error logs might be seen on FPC when reading attempt from uninitialized memory location. PR1402484
- FPC might crash after offline/online MIC-3D-16CHE1-T1-CE-H. PR1402563
- DHCP subscriber cannot reconnect over dynamic VLAN Demux interfaces because of the RPF check failure. PR1402674
- Host outbound traffic might be dropped on MPC7, MPC8, and MPC9. PR1402834

- Smg-service might become unresponsive when doing some GRE-related CLI operations. PR1403480
- The time synchronization through PTPoE might not work when an enhanced subscriber management is enabled on MX Series router. PR1404002
- Continuous kernel crashes might be observed in 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
- Incorrect output of the assigned prefixes to the subscriber in the output of the **show interface < dynamic demux interface>** command. PR1404369
- Configuration load override or load replace resets ANCP neighbors. PR1405318
- MPC might generate a core file after restarting FPC that is 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, MPC9E, and SFB2 based MX2000 routers. PR1406030
- The rpd might crash because of a race condition with the combination of community actions done at both BGP import policy and a forwarding-table policy. PR1406357
- Traffic impact might be seen if **auto-bandwidth** is configured for RSVP LSPs. PR1406822
- Layer 2 VPN might flap repeatedly after the link is up between PE and CE devices. PR1407345
- The rpd might crash when a commit check is executed on LDP trace options filtering. PR1407367
- NPC core file is created after daemon restart in jnh_get_oif_nh ( ) routine. PR1407765
- Ephemeral database might get stuck during commit. PR1407924
- Traffic forwarding fails when crossing VCF members. PR1408058
- The ToS/DSCP and TTL fields might not be copied into the outer IP header in a group VPN scenario. PR1408168
- Alarm mismatch in total memory is detected after **reboot vmhost both**. 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
- MX Series router templates are not cleaned up. PR1409398
- MX-MPC2-3D-EQ and MPC-3D-16XGE-SFPP might show Exhaust A temperature, rather than the intake temperature. PR1409406

- The 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 when the interface is disabled. PR1410465
- Kernel replication failure might be seen if an IPv6 route next hop points to an ether-over-atm-llc ATM interface. PR1411376
- Packet Forwarding Engine heap memory leak might occur during frequent flapping of PPPoE subscribers connected over aggregated Ethernet interface. PR1411389
- Virtual Route Reflector might report **DAEMON-3-JTASK_SCHED_SLIP_KEVENT** error on some hypervisor or host machine because of NTP synchronization. Routing protocol might be impacted. PR1411679
- On MX10003 routers, the 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
- Junos OS PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. PR1412659
- PPPoE subscribers might not be able to log in after unified ISSU. PR1413004
- The rpd memory leak might be seen because of an incorrect processing of a transient event. PR1413224
- 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: **[CoS(cos_rewrite_do_pre_bind_add_action:676): Binding of table 44226 to ifl 1073744636 failed, table already bound to ifl]**. PR1413297
- The support of inet6 filter attribute for the ATM interface is broken in the Junos OS Release 17.2R1 and later. PR1413663
- The services load balance might not be effective for AMS if the hash key under the **[forwarding-options]** hierarchy is configured. PR1414109
- The user might not enter configure mode because mgd is in lockf status. PR1415042
- ICMP MTU exceeded error generated from Packet Forwarding Engine does not reach the expected source. PR1415130
- The bbe-smgd process might have memory leak while running **show system subscriber-management route route-type <> routing-instance <>**. PR1415922
- After a GRES on an MX104, some tunnels might fail to pass traffic after a re-key. PR1417170
- The ECMP fast reroute protection feature might not work on MX5, MX10, MX40, MX80, and MX104. PR1417186
- An IPv4 packet with a zero checksum might not be translated to IPv6 packet properly under NAT64 scenario. PR1417215

- CGNAT with MS-MPC card does not account for AP-P out of port errors or generate a syslog message when this condition is met. PR1418128
- There is no SNMP trap message generated for **jnxHardDiskMissing/jnxHardDiskFailed** on MX10003 routers. PR1418461
- lsp-cleanup-timer is not being honored when **lsp-cleanup-timer** is configured to be greater than 2147483647. PR1418937
- Because of the PPPoE compliance issue with RFC2516, the MX Series routers allows PPPoE session-id 65535. PR1418960
- A PPP session under negotiation might be terminated if another PPPoE client bears the same session ID. PR1419500
- CPU usage on service PIC might spike while forming an IPsec tunnel under DEP/NAT-T scenario. PR1419541
- A new tunnel might not be established after changing the NAT mapping IP address until the **IPEC SA clear** command is run. PR1419542
- **rtsock_peer_unconsumed_obj_free_int: unable to remove node** from list logged extensively. PR1419647
- In the scenario where the MX Series routers and the peer device both try to bring an IPsec tunnel up, where both sides are acting as an initiator, if the peer side does not answer the MX Series ISAKMP requests, the MX Series routers can bring the peer-initiated tunnel down. PR1420293
- Failed to reload keyadmin database for **/var/etc/keyadmin.conf**. PR1421539
- VCP port reports MTU value 9152 in the ICMP MTU exceeded message while the VCP port MTU is set to 9148. PR1421629
- The remote gateway address change is not effective on MX150 platform when it is an initiator. PR1421977
- The CoS ieee-802.1 classifier might not get applied when it is configured with service activation on the underlying interface. PR1422542
- On MX10003 routers, **enhanced-hash-key symmetric** is not effective and is not shown on FPC. PR1423288
- The bbe-smgd process might crash after executing the command **show system subscriber-management route prefix <>**. PR1424054
- Interface with FEC disabled is flapping after the Routing Engine mastership switchover. PR1425211
- Soft-GRE tunnel route lost after reboot/GRES or upgrade. PR1425237
- All interfaces creation failed after NSSU. PR1425716
- Traffic loss might be seen when multiple IPsec tunnels are established with the remote peer. PR1426975
- Traffic does not flow through MACsec interfaces when configured with an unknown cipher algorithm and change back. PR1427294

- When installing YANG package without **proxy-xml** configuration statement, the CLI environment might not work. PR1427726
- L2TP subscriber and MPLS pseudowire subscriber volume accounting statistic value remains unchanged post unified ISSU. PR1429692
- The destination unreachable counter counts up without receiving any traffic. PR1431384
- The bbe-smgd process might crash if PPPoE subscribers are trying to log in when commit is in progress. PR1431459
- On MX10003 routers, PEM not present alarm raised when minimum required PEM exist in the system. PR1431926
- RSI and RSI brief should not include **show route forwarding-table** when tomcat is enabled. PR1433440
- Total number of packets mirrored after DTCP trigger add and DTCP enable is not in expected range while verifying traffic on mirror port after DTCP drop policy enable. PR1435736
- LASER TX remained enabled while the interface is disabled using the Routing Engine CLI configuration. PR1436286
- In an MPLS over UDP or MPLS over GRE scenario, if the next hop type of the MPLSoUDP/MPLSoGRE tunnel is interface route, the tunnel might not come up. PR1398362

Infrastructure

- The **show system virtual-memory | display xml validate** command displays errors. PR1356423
- SNMP OID IFOutDiscards not updated when drops increasing. PR1411303

Interfaces and Chassis

- Any filter change applied to a FTI interface triggers the FTI interface flap PR1354832
- Constant dcpfe process crash might be seen if using an unsupported GRE interface configuration. PR1369757
- The pfe_disable action does not disable the logical tunnel interfaces belonging to the affected Packet Forwarding Engine PR1380784
- Changing the value of **mac-table-size** to default might lead all FPCs to reboot. PR1386768
- When channelized interfaces are configured, the DCD core file might be generated after FPC restarts. PR1387962
- All DPCs might crash while adding or deleting a logical interface from the aggregated Ethernet bundle. PR1389206
- The dcd memory leak might be seen when committing a configuration change on static route tag. PR1391323
- Error message might be seen if GR interface is configured. PR1393676

- The dcd crash might be seen after deleting the subinterface from VPLS routing-instance and mesh-group. PR1395620
- **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 might cause interface flapping during commit. PR1402122
- The statement **targeted-broadcast** does not work on IRB interface. PR1404442
- The subscriber might not access the device because of 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 platforms. PR1410079
- OAM CFM MEP flaps might occur when hardware-assisted keepalives are enabled. PR1417707
- Monitor **ethernet loss-measurement** command returns invalid ETH-LM request for an unsupported outgoing logical interface. PR1420514
- Incorrect value on speed might cause traffic destined to the IRB's VIP to be dropped. PR1421857
- EVPN aggregated Ethernet interface flaps followed by a commit. PR1425339
- The statement **flexible-queuing-mode** is not working on MPC5E of VC member1. PR1425414
- The vrrpd process might crash after deleting VRRP sessions for several times. . PR1429906

Layer 2 Ethernet Services

- The SNMP query on LACP interface might lead to lacpd crash. PR1391545
- Log messages **dot1xd[]: task_connect: task ESP CLIENT:...: Connection refused** might be reported in Junos OS Release 17.4 or later. PR1407775
- DMAC problem of IRB interface is seen for traffic over the layer 2 circuit. PR1410970
- The IPv6 neighbor might become unreachable after the primary link goes down in VPLS scenario. PR1417209
- jdhcpd becomes aware of some of the existing configuration only after **commit full** or **jdhcpd restart**. PR1419437
- Change the nd6 next hops to reject next hops once L2 interfaces get disassociated with IPv6 entries. PR1419809
- The jdhcpd process might consistently run at 100 percent CPU and not provide service if the **delay-offer** is configured for DHCP local server. PR1419816
- jdhcpd daemon might crash during continuous stress test. PR1421569

Layer 2 Features

- The unicast traffic from IRB interface towards LSI might be dropped because of the Packet Forwarding Engine mismatching at egress processing. PR1381580
- The rpd crashes after iw0 interface is configured under a VPLS instance. PR1406472
- In a Layer 2 domain, there might be unexpected flooding of unicast traffic at every 32-40s interval towards all local CE-facing interface. PR1406807
- Broadcast traffic might be discarded in a VPLS local-switching scenario. PR1416228
- Commit error will be seen but the commit is processed when adding more than zero. PR1420082

MPLS

- The rpd process might restart after an MPLS LSP flap if **no-cspf** and **fast-reroute** are configured in an LSR ingress router. PR1368177
- The rpd might crash on backup Routing Engine after switchover. PR1382249
- An RSVP-signaled LSP might stay in down state after a link in the path flaps. PR1384929
- The rpd process might keep crashing 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
- The L2circuit information is not advertised over the LDP session if **ldp dual-transport inet-lsr-id** is different from the router-id. PR1405359
- Resources might be reserved for stale RSVP LSP when RSVP is disabled on the interface. PR1410972
- The rpd might crash if **longest-match** is configured for LDP. PR1413231
- LDP route is not present in inet6.3 if IPv6 interface address is not configured. PR1414965
- RSVP-signalled LSP takes 3 to 4 minutes before LSP switchover begins, causing the long traffic to be dropped and discarded. PR1416487
- LDP route might be missing in inet.3 when enabling sr-mapping-client on LDP-SR stitching node. PR1416516
- In OSPF multiarea topology, RSVP LSP is stuck in down state. PR1417931
- Bypass dynamic RSVP LSP tears down too soon when being used for protecting LDP LSP with the configuration statement **dynamic-rsvp-lsp**. PR1425824
- When MBB for P2MP LSP fails, it is stuck in old path. PR1429114

Network Management and Monitoring

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

Platform and Infrastructure

- MAC addresses might not be learnt on bridge domains after XE/GE interface flap. PR1275544
- The **Platform failed to bind rewrite** message might be seen when chassis control restart is done with the CoS rewrite rule configured on an aggregated Ethernet interface. PR1315437
- Move XQ_CMERROR_XR_CORRECTABLE_ECC_ERR to minor and reclassify remaining XQCHIP CMERROR from fatal to major. PR1320585
- Kernel and ksyncd core files are seen after dual cb flap at **rt_nhfind_params: rt_nhfind()** found a next hop different from that on master 30326. PR1372875
- Daemon dfwd might crash with **DFWD_TRASHED_RED_ZONE** log messages. PR1380798
- jlock hog reported at restart routing. PR1389809
- Traffic is dropped when passing through MS-DPC to MPC. PR1390541
- The lockout-period might not work for the user being locked out. PR1393839
- RVT interface might start flapping. PR1399102
- Syslog error messages: **[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_Packet Forwarding Engine:3211): Packet Forwarding Engine_id 0 cchip 0[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_Packet Forwarding Engine:3272): No Packet Forwarding Engine found for Packet Forwarding Engine_id_start 0**. PR1402377
- MAP-E some ICMP types cannot be encapsulated or decapsulated on the SI interface. PR1404239
- Some files are missing while log archiving. PR1405903
- Abnormal queue-depth counters in **show interface queue** output on interfaces that are associated to XM2 and 3. PR1406848
- IPv6 traffic might be dropped between VXLAN bridge domain and IP/MPLS network. PR1407200
- The **class-of-service** configuration changes might lead to traffic drop on cascade port in Junos fusion setup. PR1408159
- Traffic is getting dropped when there is a combination of DPC/FPC card and MPC card on egress PE router in L3VPN. PR1409523
- The VLAN tag is incorrectly inserted on the access interface if the packet is sent from an IRB interface. PR1411456
- The MPC might crash when one MIC is pulled out when this MIC is booting up. PR1414816
- Distributed multicast forwarding to the subscriber interface might not work. PR1416415
- The **op url** command cannot run a script with libraries from **/config/scripts**. PR1420976

- ARP request is not replied although **proxy-arp** is configured. PR1422148
- EX9200-12QS switch is sending tagged packets through access interface and through trunk interface with native VLAN-ID. PR1424174
- Enabling the sensor **/junos/system/linecard/qmon/** causes continuous **ppe_error_interrupt** errors. PR1434198
- BR for MAP-E does not return ICMP Type=3/Code=4 when over MTU sized packet comes with DF bit. PR1435362
- MAP-E encapsulation and decapsulation with specific parameter works incorrectly. PR1435697

Routing Policy and Firewall Filters

- On MX Series routers, the CLI statement **as-path-expand last-as** causes commit failure. PR1388159
- The rpd process might crash when **routing-options flow** configuration is removed. PR1409672

Routing Protocols

- The VRF static route might not be exported when **route-distinguisher-id** is used on route reflector in BGP L3VPN scenario. PR1341720
- Qualified next hop of static route might not be withdrawn when BFD is down. PR1367424
- The static route might persist even after its BFD session goes down. PR1385380
- BGP sessions might keep flapping on backup Routing Engine if **proxy-macip-advertisement** is configured on IRB interface for EVPN-VXLAN. PR1387720
- 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
- BGP IPv6 routes with IPv4 next hop causes rpd crash. PR1389557
- The rpd process generates a core file on backup Routing Engine during neighborship flap when using authentication key with size larger than 20 characters. PR1394082
- BGP DMZ link bandwidth is not able to aggregate bandwidth, when applying the policy. PR1398000
- The rpd soft core file might be generated when Layer 2 VPN is used. PR1398685
- The rpd process might crash in BGP setup with NSR enabled. PR1398700
- On Junos OS, the BGP packets can trigger rpd to crash when BGP tracing is enabled. PR1399141
- The UHP behavior is not supported for LDP to SR stitching scenario. PR1401214
- There might be unexpected packet drop in MoFRR scenario if active RPF path is disabled. PR1401802
- The rpd might be stuck at 100 percent when auto-export and BGP add-path are configured. PR1402140
- BGP router on the same broadcast subnet with its neighbors might cause IPv6 routing issue on the neighbor from other vendors. PR1402255
- EVPN multihoming MAC might not be installed by the remote PE device. PR1403881

- Memory leaks when labeled IS-IS transit routes 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
- On MX Series routers, the mcsnoopd process generates a core file immediately after the commit change related to VXLAN-EVPN configuration. PR1408812
- SID label operation might be performed incorrectly in OSPF SPRING environment. PR1413292
- The unexpected AS prepending action for AS path might be seen after the **no-attrset** statement is configured or deleted with **vrf-import** or **vrf-export** configuration. PR1413686
- Dynamic routing protocol flapping with VM host Routing Engine switchover on NG-RE. PR1415077
- Route information might be inconsistent between RIB and OSPF database when using OSPF LFA feature. PR1416720
- A memory leak in rpd might be seen if source packet routing is enabled for IS-IS protocol. PR1419800
- IPv6 IS-IS routes might be deleted and not be reinstalled when MTU is changed under the logical interface level for family inet6. PR1420776
- A timing issue while closing a PIM task and an auto-RP at the same time might sometimes result in generating an rpd core file. PR1426711
- The rpd might crash while handling the withdrawal of an imported VRF route. PR1427147

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
- IP ToS bits are not copied to 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- IFL** might be seen when L2TP subscribers with duplicated prefixes or framed-route log in. PR1406179
- The kmd process might crash on MX Series platforms when IKEv2 is used. PR1408974
- The jpppd core file is seen on LNS. PR1414092
- L2TP LAC might fail to tunnel static pp0 subscriber to the desired LNS. PR1416016
- IPsec SA might not come up when the local gateway address is a VIP for a VRRP configured interface. PR1422171

- In subscriber with L2TP scenario, subscribers are stuck in INIT state forever. PR1425919
- Some problems might be seen if the client negotiates LCP with no ppp-options to LAC. PR1426164

Software Installation and Upgrade

- The configuration loss and traffic loss might be seen if backup Routing Engine is zeroized and is then switched over to master Routing Engine within a short time. PR1389268
- JSU might be deactivated from FPC in case of power cycle. PR1429392

Subscriber Access Management

- Usage monitoring information AVP might activate service accounting. PR1391411
- The DHCPv6-PD client connection might be terminated after commit when RADIUS assigned address is not defined within the range of a local pool. PR1401839
- The authd crash might be seen because of a memory corruption issue. PR1402012
- JSRC used RADIUS service accounting protocol instead of JSRC for SRC installed service. PR1403835
- Continuous log message **authd[18454]: %DAEMON-3-LI: liPollTimerExpired returned 0**. PR1407923
- PPPoE session might be disconnected when LI attributes are received in access-accept with invalid data. PR1418601
- RADIUS authentication server might always be marked with DEAD. PR1429528

User Interface and Configuration

- The **show configuration** and **rollback compare** commands are causing high CPU. PR1407848

VPNs

- In a specific CE device environment in which asynchronous notification is used, after the link between the PE and CE devices goes up, the L2 circuit flaps repeatedly. PR1282875
- The receivers belonging to a routing instance might not receive multicast traffic in an extranet next-generation MVPN scenario. PR1372613
- High rpd CPU utilization on the backup Routing Engine might be observed in MVPN and NSR scenario. PR1392792
- Downstream interface is not removed from multicast route after getting PIM prune. PR1398458
- The next-generation MVPN traffic drop might be seen when **static-umh** is configured in next-generation MVPN scenario. PR1414418
- The rpd might crash in rosen MVPN scenario when a same provider tunnel source address is being used for both IPv4 and IPv6. PR1416243
- The deletion of (S,G) entry might be skipped after the PIM join timeout. PR1417344
- The rpd process might crash in rare conditions when extranet NG-MVPN is configured. PR1419891

Resolved Issues: 18.2R2

General Routing

- The **show configuration | compare** command displays the unchanged configuration after deleting a part of the configuration under the firewall section. PR1042512
- TACACS+ access does not work after an upgrade. PR1220671
- Mspmand core files are generated in rare conditions because of a high rate of TCP traffic. PR1253862
- The wrong TBB Packet Forwarding Engine component's temperature might be reported on MX80. PR1259379
- On MX Series, the **show chassis led** command should not be displayed in possible completions of the **show chassis** command. PR1268848
- An FPC crash or reboot is observed when bringing up about 12,000 Layer 2 Bitstream Access (L2BSA) subscribers simultaneously. PR1273353
- Error messages are observed on the vty session while running script for IGMP snooping over EVPN-VXLAN. PR1276947
- Error messages might be seen if the aggregated Ethernet interface hosted on MPC-3D-16XGE card flaps.PR1279607
- Migration from syslog API to Errmsg API;**/src/junos/usr.sbin/mobiled**. PR1284625
- The **apply-path** prefix is not inherited under policy after commit. PR1286987
- PPPoE cannot dial in because all the PADI packets are dropped as "unknown iif" when you deactivate or activate an aggregated Ethernet configuration. PR1291515
- Wrong packet statistics are reported in the ifHCInUcastPkts OID. PR1306656
- In a few cases, it was seen that RS are all up but virtual service is down. This was seen mainly in configuration load override conditions. PR1313009
- With autoinstallation on USB storage device is configured, interface-related commits might not take effect because of the dcd error. PR1327384
- The **PSM X Not OK** alarm is set or cleared continuously when some of the PSMs are in power-off state. PR1334572
- Tc_count counters in filter with **scale-optimized** knob, are not incrementing. PR1334580
- With certificate authority (CA), where intermediate CA profiles are not present on the device, in some corner cases, the PKI deamon 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 does not get automatically upgraded and needs to be manually upgraded after a Junos OS upgrade. PR1337028

- Link flaps or stays down because of an interoperation issue between MX Series and EX9200 device and the transport device. PR1337327
- Very few subscribers show wrong accounting values in a large-scale subscribers scenario. PR1340512
- The RLT interface might not be able to route and forward traffic in Junos OS Release 17.3. PR1344503
- The rpd might crash when the dynamic-tunnels next-hop resolving migrates to a more specific IGP route. PR1348027
- Routing Engine mastership keepalive timer is not updated after the GRES configuration is removed. PR1349049
- The MPC might crash when the MIC is removed. PR1350098
- PPE errors and async xtxn errors are observed when MPC is restarted or removed.PR1350909
- The rpd might permanently consume CPU when a logical-system configuration is committed. PR1353548
- The rpd might generates core files when adding an inter-region template routing instances. PR1354629
- The ifinfo process might crash on an MX Series device functioning as a BNG running a 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
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100 Gbps). PR1355168
- Alarm LED on the MX204 is not working to indicate the minor or major faults. PR1355225
- Packets destined to Routing Engine might be dropped in the kernel when LACP is configured. PR1355299
- VM crash might be seen when the Layer 2 circuit pseudowires are terminated. PR1355530
- The chassis alarm does not reflect the right state when the AC voltage for INP0 and INP1 is out-of-range. PR1355803
- The mpls-ipv4 templates do not have the correct source address and destination address as 4294967295, or the source mask and destination mask as 0, after the mpls-flow table size is added at runtime. PR1356118
- Executing the **show pppoe underlying-interfaces** command might cause the bbe-smgd to crash in a scaled subscriber environment. PR1356428
- Link stays up unexpectedly on the MX204 with copper cable removed. PR1356507
- DHCP subscribers fail after reconfiguration of ports from tagged to untagged mode. PR1356980
- With Junos OS Release 18.2R1, PTPoE packet exchanges do not happen with the MIC-3D-SR-4GE-2XGE MIC when PTP master and slave interfaces have **ethernet-bridge** encapsulation and are part of a bridge domain. PR1357017
- The bbe-smgd process might be stuck in subscriber scenario with External Node Slicing. PR1357252

- Upgrading from Junos OS Release 15.1F2-S20 to Junos OS Release 15.1X12 using **validate** throws the fabric mixed mode error. PR1357423
- A rpd memory leak is observed for RT_NEXTHOPS_TEMPLATE. PR1357897
- The MPC/FPC might be unable to reply to request messages of the Routing Engine in a highly scaled subscriber scenario. PR1358405
- An incorrect traffic load balanceing might be seen even if locality-bias is configured on MX Virtual Chassis. PR1358635
- The **show chassis ethernet-switch** command output on MX-TVP platforms is different from that of the MX2010 router. PR1358853
- Multiple bbe-smgd crashes might be seen when many subscribers are logging in simultaneously. PR1358868
- The **show chassis fpc** command might display **Bad Voltage** for an FPC powered off by configuration or a CLI command after the **show chassis environment fpc** command is executed. PR1358874
- The bbe-smgd restarts unexpectedly during a graceful Routing Engine switchover (GRES). PR1359290
- The FRU model number is not displayed for a few FRUs in MX10008 and MX10003 routers. PR1359300
- An IPv6 subscriber might fail to access the network. PR1359520
- A PTSP subscriber is unable to log in again after initial login attempt on PPC-based MX platforms (for example, MX104). PR1359574
- Scheduled boot for both the Routing Engines with special time format might fail. PR1359602
- The bbe-smgd might fail to add members to some of the aggregated Ethernet interfaces randomly when there are many aggregated Ethernet interfaces in the access configuration. PR1359986
- The rpd process generates core files at **Assertion failed rpd[10169]: file "../../../../../../../../src/junos/usr.sbin/rpd/lib/rt/rt_attrib.c**, line 3329: **rt_template_get_rtn_ngw(nhp) <= 1** during a Routing Engine switchover with SRTE routes. PR1360354
- An rpd scheduler slip might be seen when you frequently delete, modify, or add groups that are applied at the top level. PR1361304
- IP-over-VPLS traffic is affected by the EXP rewrite rule on the core-facing MPLS interface. PR1361429
- The MX Series device acting as a BNG does not generate the SNMP trap for ESMC/SSM quality level failed. PR1361430
- The rpd gets struck at 100 percent after the **clear bgp neighbor** operation. PR1361550
- Spontaneous bbe-smgd core file might be generated on the backup Routing Engine. PR1362188
- The MS-MPC might reset continuously on an MX Series router. PR1362271
- Traffic loss of 1 percent is seen during the GRES phase of unified ISSU from Junos OS Release 17.3-20180527.0 to Junos OS Release 17.3-20180527.0. PR1362324

- The route might get stuck after BGP neighbor and route flap. PR1362560
- Executing the **show route prefix proto ip detail** command during route churm in a scaled route scenario might lead to FPC crash. PR1362578
- Unexpected DCD_PARSE_ERROR_SCHEDULER messages are logged when the MS-MPC or MS-MIC is taken offline or brought online. PR1362734
- Rapid memory leak is seen in bbe-smgd if the service dynamic profile variable name and the associated default value are configured to be the same. PR1362810
- The nondefault routing instance is not supported correctly for NTP packets in a subscriber scenario. PR1363034
- Some CLI functions are not triggering properly (**set security ssh-known-hosts load-key-file**, **set system master-password**). PR1363475
- 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 request to record VCCP heartbeat state changes in the syslog by default. PR1363565
- Some error logs might be seen on MX2010 and MX2020 routers equipped with SFB2. PR1363587
- The xmlproxyd process for internal interfaces reports uint32 instead of uint64. PR1363766
- The Layer 2 circuit on MPC7E/8E/9E with **asynchronous-notification** and **ccc** configured might keep flapping when the circuit is going up. PR1363773
- The multicast route update might get stuck in the KRT queue and the rpd might crash if rpd and the kernel go out of synchronisation. PR1363803
- The FPM board status is missing in the SNMP MIB walk result. PR1364246
- A traffic loop might occur even though that port is blocked by RSTP in a ring topology. PR1364406
- The kernel might crash after repeatedly deactivating or activating interfaces-filter-class-of-services configurations due to accessing stale memory entry. PR1364477
- Unexpectedly large **shmlog** folder size consumes most of the disk space. PR1364775
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with the Link-Layer-Down flag set. PR1365263
- Default adapter type changed from E1000 to VMXNET3. PR1365337
- Traffic drop is seen if 3 link training failure is seen in a line card. PR1365668
- On MPC7E, the ukern crashes and the FPC reboots with the vty command **show agent sensors verbose**. PR1366249
- The MS-MPC or MS-PIC might crash in a NAT scenario. PR1366259
- On MX150, the upgrade to Junos OS Relesae 18.1R1.9 fails because installing package nfx-2-routing-data-plane-1.0-0.x86_64 needs 76 MB on the / filesystem. PR1366324

- The next hop of MPLS path might be stuck in hold state, which could cause traffic loss. PR1366562
- The SNMP MIB walk for UDP flood gives different output statistics from what the CLI gives. PR1366768
- On MX960 routers, the following syslog errors are seen **LOG : Err] Failed to allocate 2 jnh-dwords for encap-ptr(ether-da)!,LOG: Err] gen_encap_common: jnh-alloc failed! 8**. PR1366811
- Taking the fabric links of Packet Forwarding Engine 4 and Packet Forwarding Engine 5 offline is not supported. PR1367412
- The bbe-smgd process might crash during the authentication phase for an L2BSA subscriber. PR1367472
- The **show system resource-monitor fpc** command output might show a nonexisting Packet Forwarding Engine. PR1367534
- RTG interface status will be shown as incorrect status in the output of **show interface**. PR1368006
- In BBE configurations, receipt of a crafted IPv6 exception packet causes a denial of service (CVE-2018-0058). PR1368599
- The commit or commit check operation might fail due to the error of **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- Error messages about mic_sfp_phy_mdio_sgmii_lnk_op might be seen after the FPC boots up on an MX Series or EX9200 platform. PR1369382
- SNMP MIB walk causes KMD errors. PR1369938
- L2TP subscriber firewall filter might not be removed from the Packet Forwarding Engine when routing-services are enabled in the dynamic profile. PR1369968
- Kernel crash might be seen after you commit DEMUX-related configurations. PR1370015
- The rpd might crash after a Routing Engine switchover is performed or the rpd is restarted if an interface-based Dynamic GRE Tunnel is configured. PR1370174
- The packet whose size exceeds 8000 bytes might be dropped by MS-MPC in ALG scenario. PR1370582
- All the MX150 devices running VRRP on a LAN are stuck in master state. PR1371838
- The bbe-smgd process generates core files when the FPC is restarted. PR1371926
- The FPC causes high CPU utilization or crash during a hot-banking condition. PR1372193
- The smgd core files are generated after the essmd restart with reference to mmf_ensure_mapped (mmf=0xe8f0200, offset=4294967295, len=108) at ../src/junos/lib/libmmf/mmf.c:1972. PR1372223
- The Routing Engine might crash after a non-GRES switchover. PR1373079
- The AOC type optics fail to Initialize on MACsec TIC bootup. PR1373572
- URL filtering might not work when the data interfaces move from one VRF to another. PR1373582
- BOOTP packets might be dropped if BOOTP support is not enabled at the global level. PR1373807
- The cosmetic log **warning: [---] is protected, 'protocols ---' cannot be deleted** is seen after commit using **configure private** in a configuration with the **protect** flag present. PR1374244

- The FPC might be unable to work properly if one child interface is removed from an aggregated Ethernet bundle in a dynamic VLAN subscriber scenario. PR1374478
- The bbe-smgd process generates core files continiously while deleting multicast group node from the tree. PR1374530
- PCE-initiated LSPs remain in **Control status became local** state after removing the PCE configuration. PR1374596
- The rpd core files are found io_session_trace ioth_read_request_process jtask_jthr_thread_main_loop. PR1374759
- A few L2BSA subscriber logical interfaces are left behind in the SMD infra and kernel after logout. PR1375070
- SFB and PDM/PSU related information is missing in jnxBoxAnatomy MIB on high-end MX Series routers. (MX2010/2020) PR1375242
- Bbe-smgd core files might be seen after doing GRES. PR1376045
- Interface optic output power is not zero when the port has been disabled. PR1376574
- Disabling OAM might cause the Broadband Edge process to crash. PR1377090
- Packets might be dropped on the data plane in the active flow monitoring scenario. PR1377500
- After NAT64 router (with MS-MPC) translates an IPv6 fragment to IPv4 fragment, the router does not insert the right value in the identification field of the IPv4 Header. PR1378818
- ICMPv6 packets larger than 1024 might be dropped if **icmp-large-packet-check** is configured on the IDS service. PR1378852
- Traffic might get discarded when the CoS configuration is changed on a PS interface. PR1379530
- Removal of the chassisd alarms for FPCs exceeding 90 percent of power budget and exceeding 100 percent of power budget. PR1380056
- The software detects SDB STS lock deadlock and breaks the deadlock itself. The system resumes normally processing on its own. PR1380231
- The rpd might restart unexpectedly when performing GRES. PR1380298
- Encryption and decryption is not happening because the Packet Forwarding Engine discards the group-vpn member while testing, which is established using the authentication-method preshared key ASCII-text. PR1381316
- Memory leak is observed in the MS-MPC. PR1381469
- Subscribers are not able to log in after double GRES, after reboot, or after config. PR1382050
- Flows are getting exported before the active timeout. PR1382531
- MAC addresses might disappear if the interface MTU of EVPN PE is changed. PR1382966
- The configuration through NETCONF session might fail. PR1383567

- The kmd crashes with the generation of core files after the IPsec connection is brought up . PR1384205
- CoS attachment might be mistakenly removed for DHCPv4 stack when DHCPv6 stack fails to be brought up for the single session dual stack subscriber. PR1384289
- Multiple bbe-smgd core files are generated with reference to bbe_mcast_vbf_dist_policy_service_encoder( ). PR1384491
- IPsec VPN traffic might fail when passing through the MS-MPC of an MX Series router with CGNAT enabled. PR1386011
- In case an LSP is locally configured without an explicit path, the ERO object remains empty in the PCRpt generated by the PCC. PR1386935
- Uninitialized EDMEM[0x400094] Read (0x6db6db6d6db6db6d) logs seen with sampling applied to a subscriber with routing-service. PR1386948
- The pccd might crash when changing the delegation priority. PR1387419
- The bbe-smgd proces crashes and generates a core file when two DHCP subscribers with the same framed-route prefix and preference values try to log in. PR1387690
- Output of the **show class-of-service interface** command incorrectly shows adjusting application as PPPoE IA tags for DHCP subscribers. PR1387712
- The bbe-smgd might not respond to the NS message for the SLAAC client on dynamic VLAN. PR1388595
- The bbe-smgd process repeatedly generates core files and stops running as a result of long-term session database shared memory corruption. PR1388867
- IGMP group threshold exceed log message displays a wrong demux logical interface. PR1389457
- The CoS **adjustment-control-profile** configuration for application DHCP tags does not get applied. PR1390101
- There is a delay in the CLI output with second or more **show subscriber <> extensive** queries when the first session is sitting at -(more)- prompt displaying **show subscribers extensive** command output. PR1390762
- The **routing-engine-power-off-button-disable** configuration statement does not work on MX204 and MX10003 routers. PR1391548
- The bbe-smgd process might crash after commiting configuration changes. PR1391562
- MX routers serving as a DHCP server for dual-stack subscribers encounter bbe-smgd core files being generated. PR1391845
- On MX2000, fans start spinning at high speed upon inserting an FPC that was previously taken offline. PR1393256
- BBE CST telenorSweden bbe-smgd core files are generated during the inflight aggregated Ethernet reconfigure action. PR1396032

Application Layer Gateways (ALGs)

- IKEv2 negotiation might fail when IKE ESP ALG is enabled in an IKEv2 redirection scenario. PR1329611

Class of Service (CoS)

- CoS traffic control profiles might fail to be applied on an aggregated Ethernet interface in a corner scenario. PR1355498
- The 802.1P rewrite might not work on inner VLAN. PR1375189
- FPC might reboot when changing CoS mode from **hierarchical-scheduler** to **per-unit-scheduler**. PR1387987

EVPN

- The MAC entry is incorrectly programmed in the Packet Forwarding Engine, leading to some traffic being discarded. PR1231402
- MPLS label leak leads to label exhaustion and the rpd process crash. PR1333944
- The rpd might crash if the EVPN instance refers to a vrf-export policy that does not have 'then community'. PR1360437
- The l2ald memory might cross the threshold in EVPN scenario. PR1368492
- Proxy ARP might not work as expected in an EVPN environment. PR1368911
- EVPN active or the active 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
- RA packets might be sent out without using the configured virtual gateway address. PR1384574

Forwarding and Sampling

- Junos OS allows firewall filters with the same name under the **[edit firewall]** and **[edit firewall family inet]** hierarchy levels. PR1344506
- The backup Routing Engine might write dummy interface accounting records after GRES. PR1361403
- The l2ald crashes while trying the adjust **mac-table-size** configuration. PR1383665

High Availability (HA) and Resiliency

- The backup Routing Engine might go to databse prompt after performing configuration remove and restore. PR1269383
- Observed the **error: not enough space in /var on re1** error while performing unified ISSU upgrade from Junos OS Release 17.4-20180328.0 to Junos OS Release 18.2-20180416.0 . PR1354069
- Virtual Chassis-Bm cannot synchronize with Virtual Chassis-Mm when the Virtual Chassis splits and then re-forms. PR1361617

Infrastructure

- Cleanup at thread exit in the FreeBSD kernel is causing memory leaks. PR1328273

Interfaces and Chassis

- Subscribers might fail to access the device after deleting the needless aggregated Ethernet configuration. PR1322678
- The aggregated Ethernet speed changes from 1 Gbps to 10 Gbps post GRES. PR1326316
- Momentary dip in traffic is observed when GRES is performed. PR1336455
- VRRP VIP becomes unreachable after one of the interfaces is deleted. PR1352741
- **Native-vlan-id** support on ps-interface. PR1352933
- The SONET interface goes down after **keep-address-and-control** is enabled in an L2VPN scenario. PR1354713
- The aggregated Ethernet interface might flap when the link speed of the aggregated Ethernet bundle is configured to oc192. PR1355270
- The PPPoE client remains in endless loop of continuously sending IPCP configuration requests. PR1360846
- Approximately 50 percent of PPPoE subscribers (PTA and L2TP) and all ESSM subscribers are lost after unified ISSU during the DT CST stress test. PR1360870
- On all Junos OS products, the CLI allows you to configure more subinterfaces than the limit of 2048 sub-interfaces on LAG interface from Junos OS Release 17.2R1. PR1361689
- Error messages such as **ifname [ds-5/0/2:4:1] is chan ci candidate** are seen during a commit operation. PR1363536
- In rare cases, L2TP subscribers might be stuck in terminated state. PR1368650
- The EOAM LTM messages might not get forwarded after a system reboot in a CFM scenario configured with CCC interface. PR1369085
- MLPPP subscribers might be unable to negotiate sessions when the dynamic-profile name contains more than 30 characters. PR1370610
- Unified ISSU could be aborted , with the message **Timed out Waiting for protocol backup chassis master switch to complete** for an MX Series Virtual Chassis configuration. PR1371297
- The dcd process might go down when **vlan-id none** is configured for an interface. PR1374933
- The FTI IFL VNI limits are changed from (0..16777215) to (0..16777214). PR1376011
- Duplicate IP cannot be configured on both SONET (so-) interface and other interfaces . PR1377690
- Some error logs (Tx unknown LCP packet) might be reported by the bbe-smgd process on MX Series platforms. PR1378912

- The dcd is restarted unexpectedly after committing a configuration with static demux interface stacking over a ps- interface. PR1382857
- The interface-control process crashes and dcd does not restart after adding an invalid demux interface to the configuration. PR1389461

Layer 2 Ethernet Services

- STP status gets wrong after changing outer VLAN tags. PR1121564
- The MAC address might not be learned due to spanning-tree state "discarding" in kernel table after a Routing Engine switchover. PR1205373
- ZTP infra scripts are not included for MX PPC routers. PR1349249
- When DHCP subscribers are inbound (LOCAL_SERVER_STATE_WAIT_GRACE_PERIOD) state, if dhcp-service is restarted then the subscribers in this state are logged out. PR1350710
- DHCPv6 relay ignores replies from the server when renewing. PR1354212
- The DHCP lease query message is replied with incorrect source address. PR1367485
- The jdhcpd process crashes during processing of a specially crafted DHCPv6 message. PR1368377
- DHCP Relay Binding state - rebinding state counter is added to DHCPv4 and DHCPv6 binding sensors. PR1368392
- The 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
- Core files are generated at .../src/junos/usr.sbin/jdhcpd/dhcpv6/dhcpv6_option.c:1004. PR1391983

Layer 2 Features

- Addressing VPLS issues are uncovered while performing negative testings. PR1356726
- The dcPacket Forwarding Engine/fxpc process might crash on the Packet Forwarding Engines with low memory when allocating huge memory. PR1362332
- The traffic might not be transmitted correctly in a large-scale VPLS scenario. PR1371994

MPLS

- When **minimum-bandwidth** and **bandwidth** statements are present in the configuration, the bandwidth selection of the lsp is inconsistent. PR1142443
- After an MPLS LSP link flap and local repair, a new LSP instance is tried to be signaled but it might get stuck. PR1338559
- Packets destined to the master Routing Engine might be dropped in the kernel when LDP traffic statistics are polled through SNMP. PR1359956

- The Layer 2 circuit might flap after an interface goes down even if the LDP session stays up when **l2-smart-policy** is configured. PR1360255
- The rpd might crash during a P2MP LSPs churn. PR1363408
- The LSP might remain up even if no path is acceptable because of the CSPF failure. PR1365653
- The rpd process might crash after RSVP is deactivated and then re-activated globally for multiple times. PR1366243
- The rpd might crash in BGP LU and LDP scenarios. PR1366920
- RSVP authentication might fail between some Junos OS releases and cause traffic loss during local repair. PR1370182
- The next hop of static LSP for MPLS might get stuck in dead state after changing the network mask of the outgoing interface. PR1372630
- The traceroute MPLS might fail when traceroute is executed from the Juniper Networks device to other device not supporting RFC 6424. 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 the RSVP traceoptions. PR1376354
- Receipt of a specifically crafted malicious MPLS packet leads to a Junos OS kernel crash. PR1380862
- Ingress LSPs are down because of the CSPF failure. PR1385204

Network Management and Monitoring

- SNMP traps not being sent by the new master Routing Engine after a Routing Engine mastership switchover. PR1350826
- The jnxDcuStatsEntry and jnxScuStatsEntry OIDs are missing after interface configuration changes. PR1354060
- The SNMP process crashes during the polling of CFM statistics. PR1364001

Platform and Infrastructure

- MQ-chip CPQ block should report a major alarm. PR1276132
- Distributed multicast might not be forwarded to a subscriber interface. PR1277744
- Provide ability to configure the host rsyslog from the Junos OS guest. PR1341549
- Packet drop might be seen on the logical tunnel interfaces lt-x/2/x or lt-x/3/x. PR1345727
- RLT subinterfaces are not reporting statistics. PR1346403
- The lt- interface gets deleted with **tunnel-services** configuration still present. PR1350733

- When the **forwarding-class-accounting** configuration 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
- Next-hop index allocation fails and private index space is exhausted through the incoming ARP requests to management interface. PR1360039
- The **Disconnected after ISSU and before switchover** error might be seen and the FPC is restarted during a unified ISSU. PR1364514
- Authentication for adding the DTCP filter is not happening on the router and the filter is not getting added. PR1365515
- Same VLAN ID is not allowed on multiple logical interfaces of the same GR interface. PR1365640
- The qmon sensors do not work with hyper-mode enabled. PR1365990
- Subscribers over aggregated Ethernet interfaces might have tail drops, which will affect the fragmented packets due to QX-chip 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
- Traffic might drop on the newly added interfaces on the MX Series router after unified ISSU. PR1371373
- The logical tunnel interface might be unable to send out control packets generated by the Routing Engine. PR1372738
- The JNH memory leaks observed in a multicast scenario with MoFRR enabled. PR1373631
- Traffic traversing an IRB interface is not tagged with a VLAN if the packets go through an additional routing instance. PR1377526
- An FPC crash might be experienced after the FPC restarts. PR1380527
- LSI binding missing upon nd6 entry refresh after Layer 2 logical interface flap. PR1380590
- Packet drops might be seen if the packet header is over 252 bytes. PR1385585
- In Junos OS Release 18.4DCB after ifconfig goes down for a ps- logical interface, its link and admin status are not going down as expected. PR1396335

Routing Policy and Firewall Filters

- The **set metric multiplier policy** command has the potential to generate negative values, given for user-permitted inputs. PR1349462
- The rpd process might crash if **then next-hop** is configured for the LDP export policy. PR1388156

Routing Protocols

- BGP not advertising routes on newly configured Layer 3 VPN instance. PR1237006
- Multihop eBGP peering session exchanging EVPN routes can result in the generation of rpd core files when BGP updates are sent. PR1304639

- The BGP session might be stuck with high BGP OutQ value after GRES on both sides. PR1323306
- The bfd process memory leak might be observed on enabling multihop BFD session for a static route with multiple qualified next hops. PR1345041
- The rpd might generate core files while running streaming telemetry. PR1347431
- An rpd crash might be seen after a Routing Engine switchover. PR1349167
- vFPC might continuously crash on vMX platform. PR1364624
- Export policy change for BGP will trigger rpd to generate core files when OpenConfig is running. PR1366696
- Static route gets unexpectedly refreshed on commit when configured with the **resolve** option. PR1366940
- About 10 minutes traffic loss is caused by BGP flapping during unified ISSU. PR1368805
- TCP sessions might be taken down during a Routing Engine switchover. PR1371045
- Route entry might be missing when IS-IS shortcut is enabled and the MPLS link flaps. PR1372937
- The rpd might crash after the **show route detail** operational mode command is issued for the RIP route. PR1386873
- Penultimate-hop router does not install BGP LU label, causing the traffic to be discarded. PR1387746
- The FPC might crash when BGP multipath is configured with protection. PR1389379
- An rpd crash is seen when **rp-register-policy** is configured with policy with more than 511 terms. PR1394259
- Next hop of the best and the second best route have the same weight value when BGP PIC for iNet is enabled. PR1395098

Services Applications

- Selectively start ZLB Delay timer at the Packet Forwarding Engine for LAC tunnels. PR1338450
- IPsec tunnels might flap when SNMP walk is executed if IPsec is configured with DPD enabled. PR1353240
- L2TP Access Concentrator (LAC) tunnel connection request packets might be discarded on the LNS device. PR1362542
- Some subscribers might be stuck in the terminating state in an L2TP scenario. PR1363194
- The L2TP subscribers might not be able to log in successfully because of jl2tpd memory leak. PR1364774
- Accounting stop message is not sent to the RADIUS server after bringing down the L2TP subscriber. PR1368840
- IPsec-VPN IKE security-associations might get stuck in **Not Matured** state. PR1369340
- Actual-Data-Rate-Downstream might not be included in the L2TP ICRQ message. PR1370699

- NAT64 does not translate ICMPv6 Type 2 packet (packet is too big) correctly when MS-DPC is used for NAT64. PR1374255
- FTP ALG is not supported with Twice NAT . PR1383964

Software Installation and Upgrade

- Commit might fail in single-user mode. PR1368986

Subscriber Access Management

- The authd process might not be started after executing a Routing Engine switchover on the backup Routing Engine without GRES enabled. PR1368067
- Address pool does not correctly cycle to the beginning of the pool when **linked-pool-aggregation** parameter is defined. PR1374295
- RADIUS VSA's, Actual-Data-Rate-Downstream and Actual-Data-Rate-Upstream values are not compliant with RFC 4679. PR1379129
- CoA updates the subscriber with original dynamic profile if RADIUS has returned a different dynamic-profile name. PR1381230
- Some subscribers fail to get SRL service as provided in the 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
- When a subscriber manually logs out using the **clear network-access aaa subscriber username** command, the following log message is seen: **authd: gx-plus: logout: wrong state for request session-id**. PR1384599
- Multiple IPv6 IANA addresses assigned for one session in an IPv6 PD binding failure scenario. PR1384889
- Usage-Monitoring-Information AVP as part of PCRF gx-plus provisioning is causing service accounting activation. PR1391411

User Interface and Configuration

- The **max-db-size** configuration does not work on some MX Series routers. PR1363048

VPNs

- The rpd process might crash after configuration change in an Layer 2 VPN scenario. PR1351386
- In dual-homed next-generation MVPN the receipt of type 5 withdrawal removes the downstream join states for some routes. PR1368788
- MVPN source redundancy can cause possible flows outage. PR1375716

Resolved Issues: 18.2R1

Application Layer Gateways (ALGs)

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

Class of Service (CoS)

- CoS wildcard configuration is applied incorrectly after router restart. PR1325708
- The Routing Engine might get into amnesiac mode after restarting if **excess-bandwidth-share** is configured. PR1348698

EVPN

- EVPN traffic mapping to specific LSPs is not working. PR1281415
- FPC might crash if VPLS configuration is deleted. PR1324830
- Core link flap might result in an inconsistent global MAC count. PR1328956
- On deactivated ESI for packet-switched services at the physical interface (IFD) level, rpd generated a core file for EVPN-VPWS pseudowire head-end termination (PWHT). PR1332652
- On doing **restart routing**, rpd core files were generated on a provider edge (PE) router that had an EVPN-VXLAN configuration. PR1333331
- The rpd process might crash when executing CLI command **show route evpn-ethernet-tag-id**.PR1337506
- In an EVPN-VXLAN environment, BFD flaps cause VTEP flaps and cause Packet Forwarding Engine crash. PR1339084
- Traffic loss might be observed in EVPN VPWS scenario if the remote PE's interface comes down. PR1339217
- The IRB logical interface (IFL) is brought up, even if L2 interfaces are absent but IM next hops present. PR1340723
- The rpd might crash if the IRB interface and routing instance are deleted together in the same commit. PR1345519
- Traffic might be lost on Layer 2 and Layer 3 spine node in multihome EVPN scenario. PR1355165
- EVPN IRB configured with **no-gratuitous-arp-request** is still sending gratuitous ARP. PR1356360

Forwarding and Sampling

- Observing Packet Forwarding Engined core file in **Packet Forwarding Engined_process_session_state_notification_msg, Packet Forwarding Engined_timer_manager_c::remove_serv_id,Packet Forwarding Engined_delete_timer_id_by_serv_sid (serv_sid=0, serv_info=0x0) at ../../../../../../src/junos/usr.sbin/Packet Forwarding Engined/Packet Forwarding Engined_timer.cc:16** .PR1296969
- The FPC CPU might reach 100 percent constantly if shared bandwidth policer is configured. PR1320349

- Error messages about dfw_gencfg_handler might be seen during unified ISSU. PR1323795
- Some firewall filter counters might not be created in SNMP. PR1335828
- The error logical interface under VPLS might be blocked after MAC moving if the logical interfaces are on the same physical interface. PR1335880
- Commit failed when attempting to delete any demux0 unit numbers that are greater or equal to 1000,000,000. PR1348587
- With MPLS EVPN with RSVP and class-of-service-based forwarding, remote MAC is not added in the forwarding table, causing traffic to be dropped. PR1353555

General Routing

- In timing hybrid mode MX MPC2 cards are not working with ACX with vlan (native-vlan-id). PR1076666
- An rpd memory leak caused by repeated RSVP reservation state block deletes RSVP paths. PR1115686
- No warning is raised when the bridge family is configured with interface-mode trunk but without **vlan-tagging** or **flexible-vlan-tagging**. PR1154024
- Unexpected MobileNext Gateway Activation license alarm is raised when TDF gateway is configured. PR1162518
- SNMP trap sent for **PEM Input failure** alarm is not generated when single input feed fails on MX960 routers.PR1189641
- The replacement PIC might bounce when PIC PB-4OC3-4OC12-SON-SFP (4x OC-12-3 SFP) is replaced with PB-4OC3-1OC12-SON2-SFP (4x OC-3 1x OC-12 SFP) and a CLI commit is made. PR1190569
- The **Pred Fail** Fan Tray chassis alarm is renamed to **Predicted Fail**. PR1202724
- CMIC:CMIC(0/1): **Unable to deregister sub error (131072) for error(0x1b0001) for module MIC** error messages are seen on MPC5E card. PR1221337
- The error log **cc_mic_irq_status: CC_MIC(5/2) irq_status(0x1d) does not match irq_mask(0x20), enable(0x20), latch(0x1d)** is seen continuously for MIC-3D-4OC3OC12-1OC48. PR1231084
- **chassisd[9132]: LIBJSNMP_NS_LOG_NOTICE: NOTICE: netsnmp_ipc_client_connection: unix connection error: socket(-1) main_session(0x9812f80)** error messages are seen after chassis-control restart. PR1243364
- GNF sometimes resets its MPC type 9 at NSR at high scale. PR1259910
- On a vMX router, the FPC might restart unexpectedly with the message **panic (format_string=format_string@entry=0x9e509c4Thread %s attempted to %s with irq priority at %d\n**). PR1263117
- The **show chassis FPC** command does not show temperature. PR1263315
- Aggregated Ethernet incorrect counters related to PR 1261207: Incorrect counters are seen for output packets on child links for ae0 interface when configured for revertive mode. PR1273983

- For inline J-Flow, when **template-refresh-rate** and **option-referesh-rate** are configured with both packets and seconds interval options, the packets interval option is not working. PR1274206
- Software changes were provided in order to fix PR1204589 and PR1256073 that addresses the following:

  PR 1204589 - When traceroute occurs over MPLS and when the TTL expired traffic has to be generated and sent back to the source through the routing-instance, the ACX chooses the highest IP address in the routing-instance as the source, which makes it looks like the the tracepath is not correct. This behavior is modified to select the correct source address by looking into the destination routing instance and the IP address. This feature was disabled by default which has been fixed through PR1256073.

  PR1256073 - The above feature was disabled by default which is enabled with the fix of this PR. A CLI command **set system allow-6vpe-traceroute-src-select** in operational mode. PR1279191
- At commit, BSYS might log messages reporting that GNF-owned PICS do not support power-off configuration when no such configuration is present. PR1281604
- On MX Series routers with MPC7E/MPC8E/MPC9E, the threshold of corrected single-bit errors should be enhanced from 32 to 1024 and the alarm severity should changed from Major to Minor. PR1285315
- During PPPoE subscriber login, errors such as **[vbf_flow_src_lookup_enabled]** and **[failed to find iff structure, ifl]** were seen on the FPC. PR1294710
- When the system exceeds the chassis temperature limit, the log message incorrectly indicates shutdown time as 240 seconds. PR1298414
- Error messages about PEM might be seen in MX Series routers with AC PEM. PR1299284
- A chassisd core file is seen after insertion of REMX2K-X8-64 in MX2000 line routers with the older RE-S-1800x4. PR1300083
- Internal latency is high during initial subscription of sensors. PR1303393
- The mgd might crash when the Ephemeral database is used. PR1305424
- The **start shell Packet Forwarding Engine network fpc** command is not working on MX960. PR1306236
- FPC syslog errors with **Packet Forwarding Engineman_inline_ka_steering_gencfg_handler: nh not found** could mean that steering rules are not installed correctly. PR1308884
- Subscribers might not be able to access the device if dynamic VLAN is used. PR1309770
- Ninety percent of subscribers might go down after unified ISSU from Junos OS Release 16.1 to Release 17.3. PR1309983
- Utilization of **commit check** just after setting the master password can trigger improper decoding of configuration secrets. PR1310764
- The incorrect error number might be reported for syslog messages with the prefix of **%DAEMON-3-RPD_KRT_Q_RETRIES**. PR1310812
- Chassis alarm should be switched-off PEMS. PR1311574

- MX Series Virtual Chassis: BNG: IPv6 RS (router-solicit) packets are dropped in the non-default RI. This issue does not occur for the default RI. PR1313722
- The L2TP LAC might drop packets that have incorrect payload length while sending packets to the LNS. PR1315009
- The **show version detail** command gives the severity log message **mobiled: main Neither BNG LIC nor JMOBILE package is present,exit mobiled**.PR1315430
- Sensors belonging to the same producer with identical reporting intervals are not streamed in parallel. PR1315517
- The **show subscribers summary port** command does not display the correct output when subscribers are connected over pseudowire. PR1315659
- Traffic load balancing: Traffic statistics counters are not getting updated in the Junos OS Release 18.1. PR1317077
- The output from **show configuration <> | display json** might not be properly enclosed in double quotation marks. PR1317223
- Linux-based micro-kernel might panic due to concurrent update on mutable objects. PR1317961
- Adding/deleting the new Traffic Load Balancer (TLB) instance might affect other existing TLB instances. PR1318184
- CoA shaping rate is not applied successfully after unified ISSU from Junos OS Release 15.1R6.7 to Release 16.1R6.2. PR1318319
- The **show subscriber summary** displays incorrect terminated subscriber count. PR1320717
- The bbe-smgd daemon might crash after performing GRES. PR1318528
- The MPC with specific failure hardware might impact other MPCs in the same chassis. PR1319560
- Kernel core file could be seen if the number of routing instances exceeds 256. PR1319781
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX Series Virtual Chassis unified ISSU. PR1320370
- PPP inline keepalive does not work as expected when CPE aborts the subscriber session. PR1320880
- MX Series routers send the IPv6 router advertisements and the DHCPv6 advertisements before sending IPCPv6 ACK from CPE. PR1321064
- On MX Series Virtual Chassis, CoS is not applied to the Packet Forwarding Engine when a VCP link is added. PR1321184
- While running SNMP walk and with continuous server flaps for over 1 hour, for a few instances the VS summary shows as down but RS shows as up. PR1321318
- SNMP MIB walk of TLB MIB jnxTLBMIB (gives total 27,201 lines of MIB entries) for 2 TLB instances (with a total of 17 virtual services and 730 real servers) takes around 9 minutes to complete. PR1321613
- The rpd might crash when two next hops are installed with the same next-hop index. PR1322535

- The rpd might crash when the OpenConfig package is upgraded with JTI streaming data in the background. PR1322553
- MS-MIC interface logical interfaces (IFLs) remain down after many iterations of going offline and online. PR1322854
- NCP Conf-Ack/Conf-Req packets might be dropped constantly from MLPPP client on next generation broadband subscriber management . PR1323265
- The CLI commands in **show system subscriber-management route routing-instance <XXX>** hierarchy show unexpected outputs. PR1323279
- The CLI command **request vmhost halt routing-engine other** does not halt the backup Routing Engine. PR1323546
- Subscribers might fail to log in after the interface is deactivated or activated. PR1324446
- The memory leakage is seen in mosquitto-nossl daemon. PR1324531
- SNMP interface filter does not work when **interface-mib** is part of the dynamic profile. PR1324573
- SNMP values might not be increased monolithically. PR1325128
- MPC cards might drop traffic under high temperature. PR1325271
- Ping might stop working and traffic will be dropped on the channelized port if MACsec is configured on one channel. PR1325282
- IS-IS adjacency fails to establish because of packets drop on Packet Forwarding Engine. PR1325311
- MACsec session might fail to establish on MX10003 platform. PR1325331
- The VLAN demux interface does not respond to the ARP request in a subscriber scenario with MX Series routers running Junos OS Release 15.1 or later with subscriber-management enabled. PR1326450
- MACsec MKA periodic transmit interval upper limit needs to be increased. PR1326526
- On MX Series, BNG CoS service object is not deleted properly for TCP and scheduler. PR1326853
- Some of show commands were issued twice when request support information is executed. PR1327165
- Minor alarm **LCM Peer Connection un-stable** is observed on MX150 after the chassisd process starts up or restart. PR1328119
- The following message is constantly logged: **fm_feacap_sys_feature_get:Attribute DB init not yet done, reading from pvid (id: 18)** . PR1328868
- **show class-of-service interface demux0 <demux interface> Adjustment overhead-accounting mode** do not provide the expected output. PR1329212
- When an AMS bundle has a single MAM added to it, the subinterfaces do not recover after the subinterface has been disabled. PR1329498
- Host-outbound traffic is not rewriting IEEE-802.1p bits for dynamic subscriber logical Interface (IFL) Over PS interface. PR1329555

- SNMP walks for Interfaces-related MIB objects are slower than expected in a scaled configuration. PR1329931
- The **show services nat mappings address-pooling-paired** command times out and fails. PR1330207
- The alarm **Too many supplies missing in Lower/Upper zone** flaps (set/clear) every 20 seconds if a zone does not have the minimum required PSMs. PR1330720
- In a subscriber scenario, if the BGP session is created by means of the subscriber interface, then the traffic destined to the BGP advertised prefix will be dropped. PR1330737
- Rpd core files are generated on the new backup Routing Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler** after disabling NSR+GRES. PR1330750
- FPC wedge with fragmented packets occurs on LSQ interface - PT1: Head and tail out of sync. PR1330998
- A highly scaled GNF might fail to complete NSR replication after a NSR switchover. PR1331145
- Non-NEBS compliant optics might be disabled when chassis temperature exceeds non-nebs-optics-overheat-trigger. PR1331186
- The bbe-smgd process might crash after executing the command **clear ancp access-loop circuit-id <circuit id of interface set>**. PR1332096
- Inaccurate J-Flow records might be seen for the output interface and next hop. PR1332666
- On MX150 platform, **set chassis alarm management-ethernet link-down ignore** is not ignoring the alarm for FPC Mgt 0 interface. PR1332799
- Upgrading from Junos OS Release 17.3 or Release 17.4 to Junos OS Release 18.1R1 is only possible with **no-validate** command on MX10003/MX204. PR1332884
- The subinfo process might crash, and it might cause the PPPoE subscribers to get disconnected. PR1333265
- rtsblob -x prints the incorrect key. PR1333985
- AA EVPN-VXLAN causes high CPU usage on the backup rpd. PR1334235
- Two subscribers cannot reach the online state at the same time if they have an identical Frame-Route attribute value. PR1334311
- The 260G MPC with HQoS supported on Atlas (MX) went for a "restart" after unified ISSU to Junos OS Release 18.2DCB in MX2010 box. PR1334612
- MPC8E or MPC9E reports high temperature alarms and fan speed changing continuously through full and normal speed iterations. PR1334750
- The rpd crashes when performing the BGP configuration change. PR1334846
- The UID limit is reached in large-scale subscriber scenario. PR1334886
- When **show subscribers** is used and the FPC number has two digits, the interface and IPv6 address get connected together for DHCPv6 PD. PR1334904

- IPsec SA cfg name mismatch and cfg could not be pushed to the PIC. PR1334966
- Traffic drops on the MX Series LNS because of a software error or unknown family exception when traffic is destined to or coming from the MLPPP subscriber if the **routing-services** command is present in the dynamic profile used by this subscriber. PR1335276
- The master LED glows on the master and backup RCBs, while performing image upgrade on the master with GRES/NSR enabled. PR1335514
- There are hitless keychain rollover feature limitations on the MIC-MACSEC-MRATE. PR1335644
- The RIP route updates might be partially dropped when NSR is enabled. PR1335646
- The **MAC_STUCK**message might be seen on MS-MPC or MS-MIC. PR1335956
- JET application might not be respawned after a normal exit. PR1336107
- Subscriber might experience an SDB down event and drop the clients' connections when issuing **show subscribers** commands. PR1336388
- On an MX2000 with an SFB card installed, a high amount of traffic volume on MPC7E, MPC8E, or MPC9E might cause traffic drops with cell underflow messages. PR1336446
- The bbe-smgd daemon might generate core files when doing CoS configuration of logical interfaces or interface sets. PR1336852
- Configuring **lldp neighbour-port-info-display port-id** does not take effect. PR1336946
- AI-script does not get auto re-install upon a Junos upgrade on NG-Routing Engine. PR1337028
- Error log message **sdb_db_interface_remove: del ifl:si-<index> with licnese cnt non zero on** can be seen on LTS during subscriber logout. PR1337000
- FPC temperature mismatch for MPC6/8/9 occurs on MX2000 line platform. PR1339077
- DDoS counters for OSPF might not increase. PR1339364
- IPsec VPN, Session and Serviceset sensor prototype files are being added to the Junos Telemetry Interface packaging. PR1339883
- The MX10003 MPC offline button is not effective. PR1340264
- CLI shows CB states as being online after you press the RCB offline button for more than 4 seconds. PR1340431
- VRRP is stuck in the master Routing Engine during upgrade or cold boot. PR1341044
- There might be traffic loss on some subscriber sessions when more than 32,000 L2TP subscriber sessions are anchored in ASI interface. PR1341659
- The reboot of Routing Engine might occur if PPPoE interface is configured over an aggregated Ethernet or RETH interface. PR1341968
- With discard Interfaces (configured with IGMPv3), KRT queue gets stuck while deleting multicast next hop (MCNH) with error **EPERM -- Jtree walk in progress**.PR1342032

- jnxContentsType does not display details related to fixed ports and normal TIC. PR1342285
- SNMP walk might failed for LLDP related OIDs. PR1342741
- The vFPC might become absent, resulting in the total loss of traffic. PR1343170
- In an MPLS/RSVP environment, LSP might get stuck in down state with **Record route: <self> ...incomplete**.PR1343289
- Queue counters are not getting displayed in the interface details for MX150 platform once the system reboots. PR1343306
- Errors in unified ISSU because of ffp process when an upgrade from Junos OS Release 18.1 to 18.2R1 image is performed. PR1343542
- MPC7 card crashes and generates core files with DHCPv6 on static VLAN logout. PR1343965
- MX is sending IPv6 RA and the DHCPv6 advertisements before IPCPv6 Ack from CPE. PR1344472
- The ancpd process generated core files at **src/junos/usr.sbin/ancpd/ancpd_smgd.c:2299** in clearing ANCP subscribers in a scaled scenario. PR1344805
- l2cpd generates c core file (**l2cpd_ifbd_attach (ifbd=0x98914c0, vlan_id=1, line_vid=1**) after disabling mc-ae on QFX10002-60c {default vlan-scenario}, which is getting hit where Delete is missed by l2cpd because it uses sync socket read when it starts. PR1344983
- The Framed-route "0.0.0.0/0" will not be installed in MX Series platform with Junos enhanced subscriber management releases. PR1344988
- EVPN-VXLAN: ARP packet uses VRRP/virtual-gateway MAC in Ethernet header instead of IRB MAC address. PR1344990
- Rpd crash might be seen if the **no-propagate-ttl** command is set in a routing instance that has a specific route. PR1345477
- MAC address of multiple interfaces are found to be duplicate. PR1345882
- Routing Engine model changed from JNP10003-RE1 to RE-S-1600x8. PR1346054
- 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 VC-Bm when deleting one VCP link on VC-Mm. PR1346328
- The twice-napt-44 sessions are not syncing to the backup SDG with stateful sync configured. PR1347086
- IPv6 MAC resolve will fail if the DHCPv6 client uses a non-EUI64 link-local address. PR1347173
- The prerequisite of installing 32-bit libstdc++ package on host is no longer needed. PR1347921
- MIC-3D-20GE-SFP-E crashed and generated core files due to ISR 2 MIC error interrupt hogging. PR1348107

- Packet loop is detected when VRF multipath is enabled with the **equal-external-internal** command under the L3VPN instance and **install-nexthop** is enabled in the forwarding-table export policy regarding that L3VPN route. PR1348175
- Unable to set fti as output for port-mirroring instance. PR1348317
- Get-config for hidden choices is not working with ODL controller. PR1348503
- Chassisd memory leak issue occurs on MX10003 and MX204 platforms and causes eventual Routing Engine switchover and crash. PR1348753
- DHCPv6 Solicit dropped on L2TP LNS in MX Series Virtual Chassis when incoming interface is on VC-master and both anchor si- interface and VCP port on VC-backup on MPC2 NG or MPC2 NG. PR1348846
- Major alarm: **Major PEM 0 Input Failure** might be observed for DC PEM. PR1349179
- MGD crashed and generated core files due to issue in nsindb infra. PR1349288
- The MTU value for subscriber's interface might be programmed incorrectly if the command **routing-services** or **protocol pim** is configured in the dynamic profile. PR1350535
- The subinfo process might crash when executing **show subscribers address <> extensive** for a DHCPv6 address. PR1350883
- The VCP port might not come back up after it is removed and added again. PR1350845
- The Packet Forwarding Engined process is consuming 80-90 percent of CPU when running subscriber management on PPC-based routers. PR1351203
- Dynamic physical interface (IFD) creation fails when the SFP optic is plugged in MX150. PR1351387
- High CPU usage of bbe-smgd process might be seen when L2BSA subscribers get stuck. PR1351696
- After GRES, the BGP neighbors at Master Routing Engine might reset and the BGP neighbors at Backup Routing Engine take long time to establish. PR1351705
- Junos Node Slicing MSE After reinstall, one JDM server complains that the pull configuration failed and falls back to the push configuration method. PR1352503
- Bbe-smgd daemon might restart in a subscriber environment. PR1352546
- On node-sliced MX Series routers, show chassis fpc errors will not appear. PR1352705
- Offlining the MIC6-100G-CFP2 MIC using CLI command might trigger FPC card crash. PR1352921
- Rpd permanently hogging CPU due to Logical System configuration commit. PR1353548
- “3D 40x 1GE(LAN) RJ45" MIC is not recognized on MX104. PR1353632
- Syslog error: **dfw_bbe_filter_bind:1125 BBE Filter bind type 0x84 index 167806251 returned 1**. PR1354435
- Aggregated Ethernet operational state goes up even though some of the member interfaces configured under the Aggregated Ethernet are down. PR1354686

- Memory leak is found in agentd. PR1354922
- The fabric chip failure alarms are observed in GRES scenario. PR1355463
- flex-flow-sizing is not working on MX204. PR1356072
- Rpd crash was seen when issuing CLI command **show dynamic-tunnels database terse** when the system has RSVP tunnels configured. PR1356254
- I2c messages from PEM/PSM are reported if SNMP is enabled. PR1356259

High Availability (HA) and Resiliency

- The ksyncd process might crash continuously on the new backup Routing Engine after performing GRES. PR1329276
- When GRES is configured in a large-scale configuration, ksyncd crashes because of replication errors and results in insufficient available space on the hard disk. PR1332791

Infrastructure

- Cleanup at thread exit causes memory leaks. PR1328273
- The fxp0 interface not accepting IP address with **master-only** applied. PR1341325
- Junos OS is no longer going to db prompt at **~ + ctl-b**. PR1352217

Interfaces and Chassis

- IPv6 neighborship is not created on IRB interface. PR1198482
- RL-dropped packets are not displayed by **show interfaces <ifl or interface-set ifl> detail/extensive** commands. PR1249164
- L2TP subscribers might not be cleared if the access-internal routes fail to install. PR1298160
- MPC CPU usage might reach 100 percent when an OTN UFEC command is configured. PR1311154
- No route exists to the address from the directly connected route. PR1318282
- Unexpected log messages might be seen if a BGP session flaps in a dynamic-tunnels GRE scenario. PR1326983
- Unexpected log messages might be seen on a router for subscriber management. PR1328251
- Traffic loss might be seen after deleting aggregated Ethernet bundle unit 1. PR1329294
- The cfmd process crashes and generates core files.PR1329779
- The interface might not work properly after FPC restarts. PR1329896
- The dcd process might crash due to memory leak and cause commit failure. PR1331185
- The transportd process might crash when you do an SNMP query on jnxoptIfOChSinkCurrentExtTable with an unsupported interface index. PR1335438

- MX Series routers might occasionally drop the first LCP configure request packet when operating in a PPPoE subscriber management configuration. PR1338516
- When in **hardware-assited-pm-mode** and pm configuration is scale, deactivate **eth-oam** can lead to FPC crash. PR1347250
- Spontaneous jpppd generated core files on the backup Routing Engine in longevity test at **../../../../../src/junos/usr.sbin/jpppd/pppMain.cc:400**. PR1350563
- The FPC might be stuck at 100 percent for a long time when MC-AE with enhanced-convergence is configured with large-scale IFLs. PR1353397
- FPC core files related to cfmman were observed. PR1358192

Layer 2 Ethernet Services

- MX platforms might display false positive CB alarm **PMBus Device Fail**. PR1298612
- The **on-demand-address-allocation** under **dual-stack-group** does not work for IPv6. PR1327681
- The snmpget for OID dot3adInterfaceName might not work. PR1329725
- Memory leak might happen in l2cpd if the L2-learning process is disabled. PR1336720
- DHCP client is not able to connect if VLAN was modified on the aggregated Ethernet (AE) interface associated with the IRB. PR1347115
- DHCP relay agent will discard 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 crash and generate core files. PR1353597

Layer 2 Features

- An rpd process memory leak is observed upon any changes in VPLS configuration, such as deleting or re-adding VPLS interfaces. PR1335914
- VPLS instance stays in NP state after LDP session flaps. PR1354784
- RE kernel might crash when OSPFv3 is configured with IPsec key authentication over IRB interface. PR1357430

MPLS

- FPC sockets disconnects and various scheduling slips occur when executing the **show ldp traffic-statistics** command with many ECMP links and L3VPN routes. PR1214961
- The rpd might crash in LDP L2circuit scenario. PR1275766
- The **show rsvp version** command cannot display **Route Session-Id Count:** field irrespective of whether session-id is present or not. PR1285756
- Traffic drop is observed during NSR switchover for RSVP P2MP provider tunnels used by MVPN. PR1293014

- The traffic in P2MP tunnel might be lost when next-generation MVPN uses RSVP-TE. PR1299580
- The output of **show mpls container-lsp** is delayed.PR1314960
- The IPv4/IPv6 multicast traffic might get dropped in MX Series Virtual Chassis when the traffic comes in through l2circuit and goes out through aggregated Ethernet (AE) member interface across Virtual Chassis members. PR1320742
- The rpd might crash when LDP p2mp recursive is configured. PR1321626
- SNMP OID counters for mplsLspInfoAggrOctets show constant value for some LSPs even though traffic is constantly increasing in **show mpls lsp statistics** output.PR1327350
- Local repair took about 150ms > expected 100ms PR1327988
- The rpd might crash on backup Routing Engine due to memory exhaustion. PR1328974
- Fate-sharing group cost no set back to default value after CLI change, removing explicit cost configuration. PR1330161
- LDP label is generated for serial interface subnet route unexpectedly. PR1346541
- The rpd crash might happen in RSVP setup-protection scenario. PR1349036
- In a very rare scenario, rpd might crash when LDP failed to allocate self-id for the P2MP FEC. PR1349224

Network Management and Monitoring

- SNMP stops or becomes very slow after a very long period of time. PR1328455
- With **interface-mib**; MX Series router is responding with **type : NoSuchInstance** for OIDs when multiple OIDs are polled in one SNMPGET request. PR1329749
- The **show Packet Forwarding Engine statistics traffic** command output will show traffic statistics as 0 for a brief time after doing "test panic" on non-traffic-carrying line card. PR1349517
- EVENTD fails to start up with syslog configuration. PR1353364

Platform and Infrastructure

- Commit-batch is thrashing, and is not restarted. PR1284271
- DCD microbfd seems to be failing in dcd_commit_check log file even when BFD is not configured. PR1300796
- MX204 performance is degraded when using firewall filter with sampling action. PR1303529
- The source MACs might leak (or not learn) between different VPLS instances at the receiving-end VPLS PE devices. PR1306293
- VPLS instance fails to learn MAC addresses upon pseudowire switchover. PR1316459
- Rate-limit configured with small temporal buffer size might cause packet loss. PR1317385
- GNF FPC hangs at unified ISSU reboot during unified ISSU. PR1318394

- The MAC might not be learnt on MX Series routers with Trio-based card due to the negative value of the bridge MAC table limit counter. PR1327723
- The packet might get dropped in LSR if MPLS pseudowire payload does not have control word and its destination MAC starts with “4”. PR1327724
- Traffic loss might be observed on LT interface PR1328371
- The tcpdump filter might not work in egress direction on ps and lt logical interfaces (IFLs).. PR1329665
- Router hits db prompt at **netisr_process_workstream_proto**. PR1332153
- RPM MIB pingResultsMinRtt, pingResultsMaxRtt, pingResultsAverageRtt response as "1" while target address is unreachable, where it should be "0". PR1333320
- Traffic loss might be seen for some flows due to network churn. PR1335302
- Commit might fail with error reading from commit script handler: error: commit script failure. PR1335349
- The MPC might crash after setting **max-queues** to a very large number. PR1338845
- Route corruption occurs in Packet Forwarding Engine with CFM enabled on aggregated Ethernet (AE). PR1338854
- Configuring the same DHCP server in different routing instances is not supported in DHCP relay scenario. PR1342019
- Commit error occurs on configuring the same **vlan-id** on different logical interfaces (IFLs) of the same lt physical interfaces (IFDs) when **ethernet-bridge** encapsulation is configured. PR1342229
- Route corruption in Packet Forwarding Engine with connectivity-fault-management enabled for L2CKT. PR1342881
- The IPv4 GPRS traffic over aggregated Ethernet (AE) 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 continuously at 100 percent . PR1348840
- Running RSI via the console might cause system crash and reboot. PR1349332
- ICMP error messages are not generated if 'don't fragment' packets exceed the MTU of the multiservice interface. PR1349503
- 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
- Chassisd" memory leak is seen. PR1353111
- Kernel crash occurs because the initialization of logical interface (IFL) MAC filter function is missing for Packet Forwarding Engine extended port devices. PR1353498

- The FPC would crash due to the memory leak caused by the VTEP traffic. PR1356279
- Traffic is silently dropped and the following message is seen: **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH** messages. PR1357707

Routing Policy and Firewall Filters

- Access-internal route might fail to be leaked between routing instances when "from instance" is configured in the policy. PR1339689
- TPI-50840 vrf-target auto-derived internal policy is not cleaned up even after deleting the configuration, causing rpd core files.PR1357724

Routing Protocols

- The **show bgp summary** results are incorrect while assisting GR. PR1045151
- BGP extended communities with sub-type 4 are erroneously displayed at LINK_BANDWIDTH. PR1216696
- Rpd generates core files in the ASBR when BGP is deactivated in the ASBR before all stale labels have been cleaned up. PR1233893
- BGP MIBv2 enterprise MIB objects for InetAddress types not properly generate OIDs. PR1265504
- After bfdd restart seen an with next-generation MVPN and l2vpn route exchange occurs, causing MVPN and VPLS traffic drop. PR1278153
- Routing loops might be seen after configuring BGP Prefix Independent-Convergence (BGP PIC). PR1282520
- AA few adj-sid details are not updated in IS-IS database with lan + adjset scenario. PR1288331
- The lmpd will crash repeatedly when logical-system is configured on the same device. PR1294166
- MSDP sessions might flap when NSR/GRES is enabled. PR1298609
- While the device is booting up with a Junos OS 17.4R1 image, the following benign message is seen: **error: channel 0: chan_shutdown_read: shutdown() failed for fd 10 [i0 o3]: Socket is not connected during the image boot up with 17.4DCB** . PR1300409
- BGP traceoption logs are still written when deactivated. PR1307690
- Rpd core files are generated in **bgp_rt_send_message at ../../../../../../../../src/junos/usr.sbin/rpd/bgp/bgp_io.c:1460**.PR1310751
- BGP route age is getting refreshed when secondary path goes down with BGP PIC enabled. PR1312538
- The rpd might crash and generate core files with distributed IGMP. PR1314679
- The rpd might constantly consume high CPU usage in BGP setup. PR1315066
- The primary path of MPLS LSP might switch to the other address. PR1316861
- The inactive route cannot be installed in multipath next hop after disabling and enabling the next-hop interface in L3VPN scenario. PR1317623

- IS-IS might choose a suboptimal path after the metric change in ECMP links. PR1319338
- Traffic might get silently dropped temporarily when BGP GR is triggered and the direct interface flaps. PR1319631
- Issue occurs with tracing of the BGP L2VPN DF election community. PR1323596
- The rpd crash is seen when deactivating static route if the next-hop interface is type P2P. PR1323601
- When prefix limit is reached, increasing maximum-prefixes does not take effect. PR1323765
- BGP peer is not established after routing engine switchover when graceful-restart and BFD are enabled.PR1324475
- The validation replication database sometimes shows much more entries than the validation database after restarting the RPKI cache server. PR1325037
- The rpd process might crash continuously on both Routing Engines when **backup-spf-options remote-backup-calculation** is configured in IS-IS protocol. PR1326899
- Multiple next hops might not be installed for IBGP multipath route after IGP route update. PR1327904
- With BGP/LDP/IS-IS configurations, deleted IS-IS routes might still be visible in RIB. PR1329013
- The rpd might crash on backup Routing Engine after BGP peer is deleted. PR1329932
- Manual GRES with MX Series Virtual Chassis results in some packet loss on core facing interfaces. PR1329986
- The conditional route policy cannot withdraw all routes in BGP add-path scenario. PR1331615
- LDP route in inet.3 is missing when both OSPF rLFA and LFA protections are available and rejected by backup selection policy. PR1333198
- With introduction of PR1282672, discard next hop being installed when primary LSPinterface drops. When primary interface returns, discard next hop remains until BGP LU neighbor is cleared. This only impacts the cloned route (S=0). PR1333570
- Junos OS Release 15.1 onwards, IGMP joins are not processed with **passive allow-receive** configured on IGMP interface. PR1334913
- BGP sessions get stuck in active state after remote end (Cisco) restart the device. PR1335319
- Rpd core file is seen during delete and restore of BGP configuration. PR1338567
- The rpd crash might occur when receiving BGP updates. PR1341336
- Changes are required for displayed value of AIGP in **show route ... extensive** command. PR1342139
- Traffic might be silently dropped if the local device is receiving BFD-down. PR1342328
- The rpd might crash when BGP flaps. PR1342481
- 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

- Rpd might crash when BGP route damping and BGP multipath feature are configured. PR1350941
- Source community is not appended to RP (display issue in **show route** detail output) PR1353210

Services Applications

- PCP mappings cannot be manually cleared when a NAT pool is shared between PCP and standard NAT. PR1284261
- AVP 145 is not present in IRQ when ANCP DSL-type = 0. PR1313093
- SNMP MIBs are not yielding data related to sp- interfaces. PR1318339
- L2TP LTS might drop the first "CHAP Success" packet from LNS due to the delayed programming of /136 route on Packet Forwarding Engine. PR1325528
- The jl2tpd might crash if the RADIUS server returns 32 tunnel-server-endpoints. PR1328792
- Not all CSURQ messages are replied to if the number of sessions addressed in CSURQ is more than 107.PR1330150
- Crash occurs at **../src/junos/lib/libjuniper/mgmt-sock/mgmt_sock_select_info.c:35**. PR1337406
- 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 excepted. PR1348589
- While performing an SNMP walk on the IKE SAs that are getting 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 high numbers. PR1351295
- Jl2tpd process might crash shortly after one of L2TP destinations becomes unavailable. PR1352716
- L2TP tunnel-switch clients in subscriber session database reference the wrong routing-instance. PR1355396

Software Installation and Upgrade

- New versions of Junos OS do not have the tool for accessing aux port - /usr/libexec/interposer. PR1329843

Subscriber Access Management

- IP addresses are assigned discontinuously from the linked IP pools. PR1323829
- MX204 did not send **Radius Accounting-Off** message. PR1327822
- Multiple-radius-servers with different dynamic-request-port are not supported. PR1330802

- Subscriber might get stuck in terminated state when JSRC sync state is stuck in **FULL-SYNC in progress**. PR1337729
- In dual stack subscribers scenario with NDRA pool configured, the linked pools are not used when the first NDRA pool is exhausted. PR1351765

User Interface and Configuration

- CLI session might end abruptly while issuing the command **show configuration | compare rollback 1**.PR1331716

VPNs

- The rpd might crash after unified ISSU in a large-scale scenario with PIM configuration. PR1322530
- Moving MC-LAG from LDP-based pseudowire to BGP-based pseudowire might cause rpd crash. PR1325867
- The multicast might be rejected when PE devices running Junos OS received C-Mcast route from other vendors' PE devices. PR1327439
- MVPN sender-site configuration is not allowed with S-PMSI. PR1328052
- Rpd crashes and generates core files on backup Routing Engine with next-generation MVPN and NSR configuration. PR1328246
- 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
- Rpd might continuously crash on the backup Routing Engine and some protocols might flap on the master Routing Engine if 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

SEE ALSO

New and Changed Features | **112**

Changes in Behavior and Syntax | **134**

Known Behavior | **145**

Known Issues | **153**

Documentation Updates | **222**

Migration, Upgrade, and Downgrade Instructions | **223**

Product Compatibility | **230**

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 18.2R3 documentation for MX Series.

Subscriber Management Access Network

- 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

- 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 VLAN Interface

- 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 | **112**

Changes in Behavior and Syntax | **134**

Known Behavior | **145**

Known Issues | **153**

Resolved Issues | **176**

Migration, Upgrade, and Downgrade Instructions | **223**

Product Compatibility | **230**

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 Release 18.2R3, 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 18.2

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

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

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

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

Procedure to Upgrade to FreeBSD 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-x86-32-18.2R3.9-signed.tgz**

- For 64-bit Routing Engine version:

  `user@host>` **request system software add no-validate reboot *source*/junos-install-mx-x86-64-18.2R3.9-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-x86-32-18.2R3.x-limited.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-18.2R3.9-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 18.2 **jinstall** package, you cannot return to the previously installed Junos OS (FreeBSD 6.x) software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add no-validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

Procedure to Upgrade to FreeBSD 6.x based Junos OS

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

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-18.2R3.9-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-18.2R3.9-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 18.2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 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 18.2

To downgrade from Release 18.2 to another supported release, follow the procedure for upgrading, but replace the 18.2 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://pathfinder.juniper.net/feature-explorer/.

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for NFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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/software/junos/

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 18.2R3 New and Changed Features

There are no new features or enhancements to existing features in Junos OS Release 18.2R3 for NFX Series devices.

Release 18.2R2 New and Changed Features

There are no new features or enhancements to existing features in Junos OS Release 18.2R2 for NFX Series devices.

Release 18.2R1 New and Changed Features

Hardware

- **ADSL2, ADSL2+, and VDSL2 SFP modules**—Starting in Junos OS Release 18.2R1, NFX Series devices support ADSL2, ADSL2+, and VDSL2 SFP modules. The ADSL2, ADSL2+, and VDSL2 SFP modules are supported on the SFP and SFP+ ports on the NFX150 devices. Note that the ADSL2, ADSL2+, and VDSL2 SFPs are not supported on the extension modules.

  [See ADSL2 and ADSL2+ SFP Interfaces on NFX Devices and VDSL2 Interfaces on NFX150 Devices.]

Advanced Policy-Based Routing (APBR)

- **Advanced policy-based routing**—Starting in Junos OS Release 18.2R1, NFX Series devices support advanced policy-based routing (APBR), also known as application-based routing. APBR involves classifying the traffic based on the attributes of the applications and then applying filters based on these attributes to redirect the traffic. A deep packet inspection (DPI) engine is used to inspect the traffic session to identify the application. APBR provides more flexible traffic-handling capabilities by offering granular control for forwarding packets based on application attributes.

  [See Advanced Policy-Based Routing on NFX Devices.]

Security

- **Security**—Starting in Junos OS Release 18.2R1, NFX Series devices support the Layer 7 security features such as AppSecure (Application Tracking, Application QoS, Application Firewall), IPS, UserFW, and UTM.

  [See UTM User Guide for NFX Devices.]

Virtual Network Functions

- **Support for vMX VNF on NFX250-S1 and NFX250-S2**—Starting in Junos OS Release 18.2R1, vMX can be configured as a VNF on NFX250-S1 and NFX250-S2 devices. You can use the JDM CLI to configure the vMX VNF.

  [See JDM User Guide for NFX250 Network Services Platform.]

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 in Junos OS Release 18.2R3 for the NFX Series.

Factory-default Configuration

- **Plug-and-play configuration (NFX150 and NFX250 devices)**—The factory default configuration for NFX Series devices is modified to include the secure router plug-and-play configuration.
- In Junos OS Release 18.2R2, the factory-default configuration on NFX150 devices is changed to enable the front-panel copper port heth-0-3 to function as a WAN port. Previously, only the SFP ports were configured as WAN ports.

  In this release, the following changes are made to the default configuration:

  - The heth-0-3 copper port is mapped to the virtual ge-1/0/1 interface on FPC1.
  - The heth-0-4 SFP+ port is mapped to the virtual ge-0/0/3 interface on FPC0.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (NFX Series)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with a single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

SEE ALSO

- New and Changed Features | 232
- Known Behavior | 235
- Known Issues | 237

Known Behavior

IN THIS SECTION

This section lists the known limitations in hardware and software in Junos OS Release 18.2R3 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.

Hugepapges

- On NFX150 devices running Junos OS Release 18.2, 700 MB of hugepages are allocated by default for use by the system components. A portion of the 700 MB is used and the remaining free memory is available for the system as well as for third-party VNFs. PR1354027

Interfaces

- On NFX150 devices, the link will not come up if 1G SFP port is connected from heth-0-4 and heth-0-5 ports to a peer device. As a workaround, disable the auto-negotiation for the interface connected to the NFX150 device on the remote device. PR1330681
- Bi-directional 1G SFPs are not supported on 10G ports. PR1322408

- On NFX150 devices, the PPPoE session does not come up on the interface due to the hardware limitation for both tagged and untagged cases. As a workaround, enable the promiscuous mode on the interface. PR1347830

Platform and Infrastructure

- The Routing Engine boots from the secondary disk when you:
  - press the reset button on the RCB front panel, while the RE is booting up before Junos OS reboots.
  - upgrade the software by booting from the network using the request vmhost reboot network command, and the system fails to boot from the network.
  - upgrade the BIOS and it fails.
  - reboot the system and it hangs before Junos OS reboots.

  As a workaround, interrupt the boot process to select the primary disk. PR1344342

SNMP

- On NFX150 devices, SNMP does not work for the following commands:
  - **show snmp mib walk jnxIpSecTunMonOutEncryptedBytes**
  - **show snmp mib walk jnxIpSecTunMonOutEncryptedPkts**
  - **show snmp mib walk jnxIpSecTunMonInDecryptedBytes**
  - **show snmp mib walk jnxIpSecTunMonInDecryptedPkts**
  - **show snmp mib walk jnxIpSecTunMonLocalGwAddr**
  - **show snmp mib walk jnxIpSecTunMonLocalGwAddrType**

  PR1386894

Virtual Network Functions (VNFs)

- On NFX250 devices, when you configure virtual-network-functions using groups, the **show configuration** command output does not display the configurations. As a workaround, any one of the group configurations should be applied using the **set apply-groups** command. PR1302006

SEE ALSO

New and Changed Features | **232**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 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.

Interfaces

- On NFX Series devices, if the IRB interface configuration and DHCP service configuration on JDM are removed and rolled back while retaining the VLAN mapping to the IRB interface, the DHCP service fails to assign IP address to the corresponding VNF interfaces and the service chaining fails. As a workaround, remove the VLAN mapping to the IRB interface along with IRB and DHCP service configuration on JDM. PR1234055
- On NFX150 and NFX250 NextGen devices, the connectivity fault management (CFM) on CCC interface is not supported. PR1311588
- When you issue a **show interface** command on NFX150 devices to check the interface details, the system will not check whether the interface name provided is valid or invalid. The system will not generate an error message if the interface name is invalid. PR1306191
- On NFX150 devices, when you reboot the fpc0 interface, a few error messages are seen in the VTY console. PR1326487

- On NFX150 devices, the MTU of a heth interface cannot be set. The configuration knob of **set vmhost interfaces *heth-X-Y* mtu** is not supported. PR1346876
- On NFX150 devices, the link will not come up if 1G SFP port is connected from heth-0-4 and heth-0-5 ports to a peer device. As a workaround, disable the auto-negotiation for the interface connected to the NFX150 device on the remote device. PR1428020
- On NFX250-S1E devices, the JSXE0 interface might not get an IP address from DHCP server. PR1354596

Platform and Infrastructure

- On NFX150 running Junos OS Release 18.2, alternate mark inversion (AMI) does not support the ME region upgrade using the ESRT upgrade method. Hence, with the latest BIOS ABDN_U_POR3-SFP_11.37.00, BIOS upgrade for the ME region is not supported. PR1333875
- Starting in Junos Release 18.1, the file transfer rate from an external media over the network to an NFX150 device is around 40-50 Mbps. PR1290263
- Dev signed image from Juniper Networks can be upgraded in NFX150 devices. However, you can rollback if you have upgraded the dev signed image unknowingly. PR1344738
- During FTP on NFX150 devices, the following error message appears: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. PR1315605

SEE ALSO

New and Changed Features | **232**

Changes in Behavior and Syntax | **234**

Known Behavior | **235**

Resolved Issues | **239**

Documentation Updates | **240**

Migration, Upgrade, and Downgrade Instructions | **240**

Product Compatibility | **242**

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

- On NFX150 devices, the **show cli device-list** command gives a syntax error. PR1402855
- On NFX250 devcies, an SFP-T interface will not become active (UP) when it is plugged into either a ge-12/0/0 or a ge-13/0/0 interface. PR1404756
- On NFX150 devices, the **request vmhost reboot in *minutes*** command with a delay specified in minutes reboots the device immediately. PR1406018
- When you run the **show chassis fpc** or **show chassis fpc details** command, the **Temperature** field in the command output message is displayed as **Testing**. PR1433221.

Resolved Issues: 18.2R2

There are no fixed issues in Junos OS 18.2R2 for NFX Series devices.

Resolved Issues: 18.2R1

There are no fixed issues in Junos OS 18.2R1 for NFX Series devices.

SEE ALSO

New and Changed Features | 232

Changes in Behavior and Syntax | 234

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 documentation for the NFX Series devices.

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

Basic Procedure for Upgrading to Release 18.2

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

> **NOTE:** 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 Junos OS Administration Library.

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

To download and install Junos OS Release 18.2R3:

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 **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

New and Changed Features | 232

Changes in Behavior and Syntax | 234

Known Behavior | 235

Known Issues | 237

Resolved Issues | 239

Documentation Updates | 240

Product Compatibility | 242

Product Compatibility

IN THIS SECTION

- Hardware Compatibility Tool | 243
- Software Version Compatibility | 243

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 243.

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

Table 1: Software Compatibility Details with vSRX and Cloud CPE Solution *(continued)*

NFX250 Junos OS Release	vSRX	Cloud CPE Solution
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

SEE ALSO

- New and Changed Features | **232**
- Changes in Behavior and Syntax | **234**
- Known Behavior | **235**
- Known Issues | **237**
- Resolved Issues | **239**
- Documentation Updates | **240**
- Migration, Upgrade, and Downgrade Instructions | **240**

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

- New and Changed Features | **245**
- Changes in Behavior and Syntax | **255**
- Known Behavior | **261**
- Known Issues | **264**
- Resolved Issues | **269**

These release notes accompany Junos OS Release 18.2R3 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 the PTX Series.

Release 18.2R3 New and Changed Features

There are no new features or enhancements to existing features in Junos OS Release 18.2R3 for the PTX Series.

Release 18.2R2 New and Changed Features

Interfaces and Chassis

- **LACP hold-up timer configuration support on LAG interfaces (PTX Series)**—You can configure a Link Aggregation Control Protocol (LACP) hold-up timer value for link aggregation group (LAG) interfaces.

  With transport layer issues, it is possible for a link to be physically up and still cause LACP state-machine flapping, which can adversely affect traffic on the LAG interface. To prevent excessive flapping of a child (member) link of a LAG interface due to transport layer issues, a hold-up timer value is configured. LACP monitors the PDUs received on the child link for the configured time value, but does not allow the member link to transition from the expired or default state to the current state. This configuration thus prevents excessive flapping of the member link.

  To configure the LACP hold-up timer for LAG interfaces, use the **hold-time up *timer-value*** statement at the **[edit interfaces ae *aeX* aggregated-ether-options lacp]** hierarchy level.

  See [hold-time up and Configuring LACP Hold-UP Timer to Prevent Link Flapping on LAG Interfaces.]

Release 18.2R1 New and Changed Features

Hardware

- **Next-generation fixed-configuration packet transport router (PTX Series)**— Starting in Junos OS Release 18.2R1, the new PTX10002-60C features a compact, 2 U form factor that is easy to deploy in space-constrained Internet exchange locations, remote central offices, and embedded peering points throughout the network, including cloud-hosted services. The PTX10002-60C has 60 QSFP28 ports that you can configure as 100 Gbps or 40 Gbps interfaces or channelize as four 10-Gbps interfaces. The ports handle up to 6 Tbps of throughput and 4 Bpps of forwarding capacity. The PTX10002-60C is available with either AC or DC power supplies, and it has airflow out, where air comes into the vents in the port panel and exhausts through the field-replaceable unit (FRU) panel.
- **PTX10K-LC1105 MACsec line card on PTX10008 and PTX10016 routers**—Starting in Junos OS Release 18.2R1, the PTX10K-LC1105 line card provides thirty 100-Gbps or 40-Gbps QSFP28 ports with MACsec features.

  [See PTX10000 Line Card Components and Descriptions.]

Class of Service (CoS)

- **Support for class of service (CoS) on PTX10002-60C routers**—Starting in Junos OS Release 18.2R1, PTX10002-60C routers support CoS.

  CoS is the assignment of traffic flows to different service levels. Service providers can use router-based CoS features to define service levels that provide different delay, jitter (delay variation), and packet loss characteristics to particular applications served by specific traffic flows.

  [See Understanding CoS CLI Configuration Statements on PTX Series Routers.]

High Availability (HA) and Resiliency

- **Resiliency support for PTX10008 and PTX10016 routers with JNP10K-RE1**—Starting with Junos OS Release 18.2R1, resiliency support is enabled for PTX10008 and PTX10016 routers with the JNP10K-RE1 Routing and Control Boards (RCBs).

Interfaces and Chassis

- **Support for PTX10K-LC1105 line card (PTX10008)**—Starting with Junos OS Release 18.2R1, PTX10008 routers support the PTX10K-LC1105 line card. The line card is designed to provide secure Ethernet communication across high-speed links. The card consists of 30 QSFP+ or QSFP28 Pluggable ports that are Media Access Control Security (MACsec) capable. The ports support speeds of 100 Gbps or 40 Gbps, which can be configured using the CLI.
- **Protection against distributed denial-of-service (DDoS) attacks (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, PTX10002-60C devices support DDoS protection for many Layer 2 and Layer 3 protocol families and packet types. DDoS attacks typically use network control packets to trigger a large number of exceptions in the network, consuming resources and crippling network operations. DDoS protection uses firewall filters and policers available in Junos OS to discard or rate-limit control plane traffic so that malicious traffic does not overwhelm and bring down a device. To configure DDoS protection, use the **ddos-protection** statement at the **[edit system]** hierarchy level to specify the desired protocol groups, control packet types, and filter parameters.

  [See Understanding Distributed Denial-of-Service Protection on PTX Series and QFX Series Devices.]

- **Channelization support (PTX Series)**—Starting with Junos OS Release 18.2R1, you can use channelization functionality to subdivide a larger flexible optical interface into subinterfaces or channels. PTX Series routers have 12 ASIC circuits (PE) as a part of a Packet Forwarding Engine, and each PE switch has 5 ports (one standalone MAC port and 4 channelized MAC ports). The standalone MAC ports cannot be channelized. On the router, you can channelize 48 ports out of the available 60 ports.

  By default, the ports come up in a mode that does not support channelization.

  To enable channelization on an interface:

  ```
  [edit chassis fpc fpc-slot pic pic-slot]
  user@switch# set port port-number speed speed
  ```

[See Channelizing Interfaces.]

- **Enhanced fault management features**—Starting with Junos OS Release 18.2R1, PTX10001 routers support the configuration of error thresholds and actions at the error scope and error category levels. Use the **set chassis fpc *fpc-slot* error scope *error-scope* category *category* (fatal | major | minor) threshold *error-threshold* action (alarm | disable-pfe | get-state | offline | log | reset)** command to configure an error threshold and action for a particular error scope and category at the FPC level. This feature can also be configured at the chassis level (at the **[edit chassis]** hierarchy). You can also disable an error or modify the severity of a particular error at the error ID level.

  You can use the **show chassis fpc errors** command to view the error information at the error scope and category level.

Junos Telemetry Interface

- **Streaming OpenConfig data from Routing Engine sensors over UDP in protobuf format (MX Series, PTX Series, QFX Series)**—Starting in Junos OS Release 18.2R1, you can stream OpenConfig-based sensor data from Routing Engine sensors by using the Junos Telemetry Interface (JTI). JTI enables to stream the OpenConfig sensor data in gRPC/protobuf format rather than in key/value pairs. Using the protobuf format is more efficient and makes the messages smaller.

  [See Overview of the Junos Telemetry Interface.]

- **Routing Engine state sensors for the Junos Telemetry Interface (MX Series, PTX Series)**—Starting with Junos OS Release 18.2R1, you can export statistics for the Routing Engine state through the Junos Telemetry Interface using the following resource paths:

  - **/junos/kernel-ifstate/stats/churn-rate**
  - **/junos/kernel-ifstate/stats/peer-consumption-rate**
  - **/junos/kernel-ifstate/stats/vetos-statistics**

  Only gRPC streaming is supported.

  To provision the sensor to export data through 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.

  Support for the Junos Telemetry Interface was introduced on QFX10000 and QFX5200 switches in Junos OS Release 17.2R1.

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

- **Expanded ON_CHANGE support for Junos Telemetry Interface (JTI) (MX Series, PTX Series)**—Starting with Junos OS Release 18.2R1, OpenConfig support through remote procedure call (RPC) and JTI is extended to support additional ON_CHANGE sensors for some endpoints under resource paths **/interfaces/interface/state** and **/interfaces/interface/subinterfaces/subinterface/state/**.

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.

To provision a sensor to export data through 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.

To enable ON_CHANGE support, configure the sample frequency in the subscription as zero. When you create a subscription using a top-level container as the resource path (for example, **/interface**), leaf devices under the resource path **/interface** with ON_CHANGE support are automatically streamed based on events. Other leaf devices will not be streamed.

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.

[See Understanding OpenConfig and gRPC on Junos Telemetry Interface and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **J-Insight Device Monitor (PTX Series)**—J-Insight is a data-driven device-monitoring solution that provides visibility and insight into the health of a running system. Starting with Junos OS Release 18.2R1, the J-Insight framework facilitates real-time monitoring of system resources for FPC FRUs. It also has been integrated with the existing connectivity error management infrastructure to normalize error detection, monitoring, and reporting. J-Insight is an on-premise system application that uses the Junos Telemetry Interface to continuously collect data that reflect the current state and health of the device component being monitored.

  [See J-Insight Device Monitor Overview.]

Layer 3 Features

- **Support for Layer 3 unicast features on PTX10002-60C**—Starting in Junos OS Release 18.2R1, PTX10002-60C routers support the following Layer 3 features for unicast IPv4 and IPv6 traffic:
  - OSPF
  - IS-IS
  - BGP

MPLS

- **LDP support (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, the PTX10002-60C router supports the Label Distribution Protocol (LDP). LDP is a protocol for distributing labels in non-traffic-engineered applications. LDP enables routers to establish label-switched paths (LSPs) through a network by mapping network-layer routing information directly to data link layer-switched paths. [See MPLS Applications User Guide for Routing Devices.]
- **RSVP support (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, the PTX10002-60C router supports RSVP. RSVP is a resource reservation setup protocol that is used by both network hosts and

routers. Hosts use RSVP to request a specific class of service (CoS) from the network for particular application flows. Routers use RSVP to deliver CoS requests to all routers along the datapath. RSVP can also maintain and refresh states for a requested CoS application flow. [See MPLS Applications User Guide for Routing Devices.]

- **MPLS capabilities (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, MPLS capabilities are available on the PTX10002-60C router. MPLS provides both label edge router (LER) and label-switching router (LSR) capabilities, and supports the following features:
  - Object access method, including ping, traceroute, and Bidirectional Forwarding Detection (BFD)
  - Fast reroute (FRR) which is a component of MPLS local protection. Both one-to-one local protection and many-to-one local protection are supported.
  - Loop-free alternate (LFA)
  - IPv6 Provider Edge (6PE) and IPv6 VPN Provider Edge (6VPE) devices
  - Layer 3 VPNs for both IPv4 and IPv6
  - Layer 2 circuit

  [See MPLS Applications User Guide for Routing Devices.]

- **Support for IS-IS segment routing (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, IS-IS segment routing support is enabled through MPLS. Junos OS IS-IS implementation allocates node segment label blocks to support segment routing node segments. It also provides a mechanism to the network operator to provision an IPv4 or IPv6 address family node segment index. To configure segment routing (also known as source packet routing), use the following configuration statements at the **[edit protocols isis]** hierarchy level:
  - **source-packet-routing**
  - **node-segment**
  - **use-source-packet-routing**
  - **no-advertise-adjacency-segment**

  [See IS-IS User Guide.]

- **Egress peer engineering of service labels (such as BGP and MPLS) and egress peer protection for BGP-LU (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, you can enable traffic engineering of service traffic, such as MPLS LSP traffic between autonomous systems (ASs), by using BGP-labeled unicast for optimum utilization of the advertised egress routes. You can specify one or more backup devices for the primary egress AS boundary router. Junos OS installs the backup path in addition to the primary path in the MPLS forwarding table, which enables MPLS fast reroute (FRR) when the primary link fails. It provides support for the FRR protection backup scheme to perform an IP lookup to determine a new egress interface.

  [See Configuring Egress Peer Traffic Engineering by Using BGP Labeled Unicast and Enabling MPLS Fast Reroute.]

- **IPv6 tunneling over an MPLS-based IPv4 network (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, tunneling enables you to connect IPv6 sites over an IPv4 MPLS-enabled backbone. IPv6 packets are carried over an IPv4 MPLS tunnel. To enable this service, you need to deploy provider edge (PE) routers that can run IPv4, MPLS, and BGP toward the core and IPv6 toward the edge.

  [See Example: Tunneling IPv6 Traffic over MPLS IPv4 Networks.]

- **MPLS inter-AS link protection (PTX10002-60C)**—Starting in Junos OS Release 18.2R1, MPLS inter-AS link protection is supported. Link protection is essential in an MPLS network to ensure traffic restoration in case of an interface failure. The ingress router will then choose an alternate link through another interface to send traffic to its destination.

  For an MPLS inter-AS environment, link protection can be enabled when **labeled-unicast** is used to send traffic between autonomous systems (ASs). To configure link protection on an interface, the **protection** statement is introduced at the **[edit protocols bgp group *group-name* family inet labeled-unicast]** hierarchy level.

  [See Understanding MPLS Inter-AS Link Protection.]

Multicast

- **Support for multicast protocols (PTX10002-60C) routers**—Starting in Junos OS Release 18.2R1, PTX10002-60C routers support the following multicast protocols:

  - Protocol Independent Multicast sparse mode— PIM sparse mode enables efficient routing to multicast groups with receivers sparsely spread over multiple networks. To configure PIM sparse mode, include the **pim** statement at the **[edit protocols]** hierarchy level. PIM sparse mode supports static RP addresses, bootstrap routers, automatic RP announcement and discovery, and anycast RP functionality.

    [See Understanding PIM Sparse Mode.]

  - PIM source-specific multicast (PIM SSM)— PIM source-specific multicast uses a subset of PIM sparse mode and IGMPv3 to enable a client to receive multicast traffic directly from the source. PIM source-specific multicast uses the PIM sparse-mode functionality to create a shortest-path tree (SPT) between the client and the source, but builds the SPT without the help of a rendezvous point.

    [See Understanding PIM Source-Specific Mode.]

  - Internet Group Management Protocol (IGMP)—IGMP manages the membership of hosts and routing devices in multicast groups.

Network Management and Monitoring

- **sFlow functionality introduced on PTX1000 and PTX10000**—Starting in Junos OS Release 18.2R1, the PTX1000 and PTX10000 routers support sFlow, a network monitoring protocol for high-speed networks. With sFlow, you can continuously monitor tens of thousands of ports simultaneously. The mechanism used by sFlow is simple, not resource intensive, and accurate. An sFlow agent embedded in a network device samples packets and gathers interface statistics and sends the information to a monitoring station called a *collector* for analysis. An sFlow agent can be implemented in a distributed model. In such a case,

each subagent has a separate subagent ID and is responsible for monitoring a set of network ports. The subagents share a common agent address.

[See Configuring sFlow Technology for Network Monitoring (CLI Procedure) and sflow.]

- **Support for Junos Space Service Now (PTX10008 and PTX10016)**—Starting in Junos OS Release 18.2R1, PTX10008 and PTX10016 routers support Junos Space Service Now. Junos Space Service Now is an application that runs on the Junos Space Network Management Platform to automate fault management and accelerate issue resolution.

  [See Junos Space Service Now.]

- **Support for port mirroring on PTX10002-60C routers**—Starting in Junos OS Release 18.2R1, PTX10002-60C routers supports port mirroring. Port mirroring copies packets entering or exiting a port and sends the copies to a local interface for local monitoring. You can use port mirroring to send traffic to applications that analyze traffic for purposes such as monitoring compliance, enforcing policies, detecting intrusions, monitoring and predicting traffic patterns, correlating events, and so on.

  [See Configuring Port Mirroring.]

Operation, Administration, and Maintenance (OAM)

- **Connectivity fault management (CFM) support (PTX Series)**—Starting with Junos OS Release 18.2R1, PTX5000 routers with FPC-P2 support Ethernet OAM CFM on the child links of tagged aggregated Ethernet bundles for IPv4 traffic, thereby enabling you to monitor faults on those child links.

  The CFM supports fault monitoring and path discovery functionalities.

  > **NOTE:** To enable CFM on an Ethernet interface, you must configure maintenance domains, maintenance associations, and maintenance association end points (MEPs).
  >
  > [See IEEE 802.1ag OAM Connectivity Fault Management Overview .]

Routing Policy and Firewall Filters

- **Support for firewall filters and policers on PTX10002-60C routers**—Starting in Junos OS Release 18.2R1, you can define firewall filters on the PTX10002-60C routers that define whether to accept or discard packets. The PTX10002-60C routers support IPv4 filters, IPv6 filters, and MPLS filters.

  You can also use policing to apply limits to traffic flow and specify the action to be taken for packets that exceed those limits.

  [See Firewall Filters Overview.]

Services Applications

- **Support for multiple flow collectors for inline flow monitoring (PTX Series)**—Starting in Junos OS Release 18.2R1, you can export flow records generated by inline flow monitoring to four collectors under a family with the same source IP address. The Packet Forwarding Engine can export the flow record, flow record

template, option data, and, option data template packet to all configured collectors. You can configure the multiple collectors at the **[edit forwarding-options sampling instance *instance name*]** hierarchy level.

[See Monitoring Network Traffic Flow Using Inline Flow Monitoring on PTX Series Routers.]

- **Support for inline flow monitoring (PTX10008 and PTX10016)**—Starting in Junos OS Release 18.2R1, Junos OS supports inline active flow monitoring. Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Version 9 template is supported for IPv4, IPv6, and MPLS. IPFIX template is supported for IPv4, IPv6, and MPLS. Both IPFIX and version 9 templates use UDP as the transport protocol.

  [See Monitoring Network Traffic Flow Using Inline Flow Monitoring on PTX Series Routers.]

- **Support for MPLS, MPLS-IPv4, and MPLS-IPv6 inline active flow monitoring (PTX Series)**—Starting in Junos OS Release 18.2R1 on PTX Series routers, you can perform inline flow monitoring for MPLS, MPLS-IPv4, and MPLS-IPv6 traffic. Both IPFIX and version 9 templates are supported. Inline flow monitoring for MPLS-over-UDP flows was supported in Junos OS Release 18.1R1.

  [See Configuring Inline Active Flow Monitoring on PTX Series Routers.]

Software Installation and Upgrade

- **Zero Touch Provisioning (PTX3000, PTX5000, PTX10008, PTX10016)**—Starting in Junos OS Release 18.2R1, Zero Touch Provisioning (ZTP) is supported to automate the provisioning of the device configuration and software image with minimal manual intervention.

  When you physically connect a router to the network and boot it with a factory configuration, the router attempts to upgrade the Junos OS software image automatically and autoinstall a configuration file from the network through the management interface on PTX5000, PTX3000, PTX10008, and PTX10016 routers. The router uses information configured on a DHCP server to locate the necessary software image and configuration files on the network. If you have not configured the DHCP server to provide this information, the router boots with the preinstalled software and factory-default configuration. The ZTP process either upgrades or downgrades the Junos OS version.

  [See Understanding Zero Touch Provisioning and Configuring Zero Touch Provisioning.]

- **ZTP support (PTX10002-60C switch)**—Starting with Junos OS Release 18.2R1, Zero Touch Provisioning, automates the provisioning of the device configuration and software image with minimal manual intervention, and is supported on PTX10002-60C VM hosts. When you physically connect a supported device to the network and boot it with a factory configuration, the device attempts to upgrade the Junos OS software image automatically and autoinstall a configuration provided on the DHCP server.

  [See Zero Touch Provisioning.]

System Management

- **Support for request vmhost and show vmhost commands (PTX10002-60C switches)**—Starting in Junos OS Release 18.2R1, many of the **request system** and **show system** commands have been replaced with **request vmhost** and **show vmhost** commands.

Here is a list of the vmhost commands that are now supported:

- request vmhost cleanup
- request vmhost file-copy
- request vmhost halt
- request vmhost hard-disk-test
- request vmhost power-off
- request vmhost power-on
- request vmhost reboot
- request vmhost snapshot
- request vmhost software add
- request vmhost software rollback
- request vmhost zeroize
- show vmhost bridge
- show vmhost crash
- show vmhost hard-disk-test
- show vmhost hardware
- show vmhost information
- show vmhost logs
- show vmhost management-if
- show vmhost netstat
- show vmhost processes
- show vmhost resource-usage
- show vmhost snapshot
- show vmhost status
- show vmhost uptime
- show vmhost version

[See VM Host Operations and Management for more information.]

- **New tool to detect high CPU utilization for routing protocol process (PTX Series)**—Starting in Junos OS Release 18.2R1, a flight recorder tool is introduced to gather historical data on when the CPU utilization for the routing protocol process on a device was high and what processes caused the high utilization. The tool collects snapshots of data, enabling detection of high CPU usage and faster resolution of issues.

Because some of the high CPU utilization cases are intentional or expected, you can enable and disable the flight recorder tool to avoid false alarms.

[See request flight-recorder set high-cpu and show flight-recorder status.]

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 in Junos OS Release 18.2R3 for the PTX Series.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (PTX Series)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option **commit fast-synchronize** is disabled from the CLI.

Interfaces and Chassis

- **Power supply alarm is not raised when the input switch status is OFF or power is not connected (PTX10008, PTX10016)**—Starting in Junos OS Release 18.2R1, the power supply alarm **A power supply input has failed** is not raised if the INP1/INP2 switch status if off and the power is not connected. In earlier releases, an alarm is raised for the power entry module (PEM) that are not powered on as **Not Powered** irrespective of the switch state. Now, to know the power supply status, execute the **show chassis power** or **show chassis power detail** CLI command. The **DC input** is the new output parameter that provides information about the status of the input feed.

  Previous behavior:

  user@host> show chassis power

```
PEM 0:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC output: 864 W (zone 0, 72 A at 12 V, 34% of capacity)

PEM 1:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC output: 864 W (zone 0, 72 A at 12 V, 34% of capacity)

System:
  Zone 0:
      Capacity:          7500 W (maximum 7500 W)
      Allocated power:   6525 W (975 W remaining)
      Actual usage:      2616 W
  Total system capacity: 7500 W (maximum 7500 W)
  Total remaining power: 975 W

...
```

  Current behavior:

user@host> show chassis power

```
PEM 0:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (No feed expected, Both feed connected)
  DC output: 576 W (zone 0, 48 A at 12 V, 23% of capacity)

PEM 1:
  State:     Online
  Capacity:  2500 W (maximum 2500 W)
  DC input:  OK (No feed expected, Both feed connected)
  DC output: 576 W (zone 0, 48 A at 12 V, 23% of capacity)

...
```

[See show chassis power.]

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (PTX Series)**—In Junos OS Release 18.2R3, 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 that an interface holds before it changes state from down to up. In earlier Junos OS releases, the LACP hold-up information for all interfaces were displayed in a single **<lacp-hold-up-information>** XML tag. Now, for each interface, the LACP hold-up information is displayed in a separate **<lacp-hold-up-information>** XML tag.
- **New option for configuring IP addresses when the Routing Engine is the current master**—In Junos OS Release 18.2R3, a new option, **master-only**, is supported on routers with RE-MX-X6, RE-MX-X8, and RE-PTX-X8 Routing Engines at the following hierarchies:
  - **[edit vmhost interfaces management-if interface (0|1) family inet address *IPv4 address*]**
  - **[edit vmhost interfaces management-if interface (0|1) family inet6 address *IPv6 address*]**

  In routing platforms with dual Routing Engines and VM host support, the **master-only** option enables you to configure the IP address to be used for the VM host when the Routing Engine is the current master. The master Routing Engine and the backup Routing Engine can have independent host IP addresses configured. In earlier Junos OS releases, the same IP address is applied on the master and backup Routing Engines, resulting in configuration issues.

Junos OS XML API and Scripting

- **Junos XML protocol <open-configuration> operation no longer emits an uncommitted changes warning (PTX Series)**—Starting in Junos OS Release 18.2R1, the Junos XML protocol **<open-configuration>** operation does not emit an "**uncommitted changes will be discarded on exit**" warning message when

opening a private copy of the candidate configuration. However, Junos OS still discards the uncommitted changes upon closing the private copy.

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (PTX Series)**—Starting in Junos OS Release 18.2R2, 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.

Junos Telemetry Interface

- **Change to the configuration location for gRPC-based sensor subscriptions from an external collector (PTX Series)**—Starting in Junos OS Release 18.2R1, when an external streaming server, or collector, provisions sensors to export data through gRPC on devices running Junos OS, the sensor configuration is committed to the **junos-analytics** instance of the ephemeral configuration database, and the configuration can be viewed by using the **show ephemeral-configuration instance junos-analytics** operational command. In earlier releases, the sensor configuration is committed to the default instance of the ephemeral configuration database.

MPLS

- **New debug statistics counter (PTX Series)**—The **show system statistics mpls** command has a new output field, called **Packets dropped, over p2mp composite nexthop**, to record the packet drops over composite point-to-multipoint next hops.

Network Management and Monitoring

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a nondefault routing instance and nondefault logical system (PTX Series)**—In Junos OS Release 18.2R1, 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.]

- **Junos OS does not support management of YANG packages in configuration mode (PTX Series)**—Starting in Junos OS Release 18.2R2, adding, deleting, or updating YANG packages using the **run** command in configuration mode is not supported.

- **No chassis alarm when power consumption by an FPC exceeds 90% or 100% of the allocated power budget**—Starting in Junos OS Release 18.2R2, the PTX5000 routers do not raise a chassis alarm in the following events:

  - Power consumption by an FPC exceeds 90% of the allocated power budget.

  - Power consumption by an FPC exceeds 100% of the allocated power budget (in this case, a system log is registered).

- **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 18.2R2, 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 error severity (PTX10016)**—Starting in Junos OS Release 18.2R3, 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). If this error occurs, the message is displayed for informational purposes only. 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]

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 Junos OS Release 18.2R3, under rare circumstances, the host interface might 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. To avoid this issue, 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

- **New DHCP option introduced for ZTP retry (PTX Series)**—Starting in Junos OS Release 18.2R1, a new DHCP option is introduced to set the timeout value for the file downloads over FTP. If the **transfer-mode** is set as FTP, the default value for the timeout is automatically set as 120 minutes. That is, if the FTP session gets interrupted due to loss of connectivity in the middle of a file transfer, it will time out after 120 minutes and ZTP will attempt to retry the file-fetching process. This value can be overridden using the DHCP option as follows:

```
option NEW_OP.ftp-timeout code 7 = text;
option NEW_OP.ftp-timeout “val”;
```

where **“val”** is the user configurable timeout value in seconds and must be provided within double quotation marks (for example, "val").

- **ssh-keygen output is tagged in XML (PTX1000)**—In Junos OS Release 18.2R2, the ouput of the ssh-keygen utility that is invoked when generating the ssh keys, is now in its XML form, and is wrapped in **<output>** tags. You can see this 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 IA renew requests (PTX Series)**—Starting in Junos OS Release 18.2R2, 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

Known Behavior

IN THIS SECTION

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 for the PTX Series.

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

General Routing

- When an FPC goes offline or restarts, FPC x sends traffic to FPC y. The following error messages are displayed and a corresponding alarm is set on the destination FPC. Specific to PTX10000 line of devices, the transient alarm is set when this condition occurs. The alarm clears later when 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 error** PR1268678
- In the specific case of semigraceful RCB reboot initiated by the internal shell command: **vhclient init 0**, GRES takes longer to complete; that is, 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

- MPLS Ingress LSP statistics are not supported. PR1337814
- When unsupported sensors are configured, the sensors are subscribed to on the device, but no data is exported. PR1339559
- Micro-BFD configuration with interface addresses is not supported on PTX Series devices (and QFX Series devices) on FPC3. PR1341513
- The Routing Engine boots from the secondary disk when you:

  a) Press the reset button on the RCB front panel while the Routing Engine is booting up but before Junos OS is up.

  b) Upgrade software, by booting from the network using the **request vmhost reboot network** command, and the system fails to boot from the network.

  c) Upgrade BIOS and the upgrade fails.

  d) Reboot and the system hangs before Junos OS is up. PR1344342
- Due to a ZH ASIC limitation, MAC statistics under **show interface** in the Routing Engine might not reflect **Mac error Counters** properly if the ingress packet size is greater than the default mtu (1518) or user configured mtu size (**set interface <interface-name>mtu <288...9600>**PR1345779
- Ingress LSP statistics is not supported. Only transit LSP statistics are supported, and it is limited to 24000 only. PR1355909
- The 100-Gigabit DAC connected between QFX5200 and PTX10002-60C or QFX10002-60C does not link up. This is because BCM-based devices have link-training enabled while PE-based devices do not have link-training enabled for 100G DAC/CR4. PR1356834
- Frequent speed changes on interface ports might cause the the relevant port's physical interface not to be created. PR1367946
- The ingress interface and mirror interface should have the same MTU, or the mirror interface should have a higher MTU than the ingress interface. PR1372321
- The **set interfaces <interface-name>gigether-options fec <fec74/fec91/none>** configuration is not supported for JNP hardware running Junos for PTX1000-M20C. PR1388140
- A LU-BGP traffic loss is seen on link-up. When link is made up, there might be a race condition between ingress using the new path and transit programming the label route. Because, the ingress might prematurely use the new link that is dropped at transit as transit has not programmed the label route yet. A v4/v6 traffic loss is seen on link up/down. In some scenario, indirect next-hop changes from old path to better new path. But the indirect next-hop starts using the new better path before the new better path (forwarding next-hop) gets programmed in Packet Forwarding Engine. A v4/v6 traffic loss is seen on link up. When link is made up, there might be a race condition between ingress using the new path by doing indirect next-hop change and transit doing indirect next-hop change to use a better path. If ingress is faster than transit in doing indirect next-hop change, then, the traffic at transit is forwarded

to ingress and ingress loops it back to transit resulting in a micro loop until transit programs the indirect next-hop change. PR1400784

- When a member link gets deleted or deactivated from an aggregated Ethernet bundle or the the link goes down, on which the inline BFD session is currently established, the BFD session might flap. This is a day-one limitation of the inline BFD design on PTX Series routers. PR1401342

Infrastructure

- When Layer 3 interface comes up, there can be mismatch in IFL counters between Routing Engine and Jvision. This mismatch pertains to ARP/GARP packets. As ARP/GARP packet that gets initiated the moment Layer 3 interface comes up (from spirent/DUT) Routing Engine ends up having one packet less on IFL. PR1361282

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
- On PTX10008 and PTX10016 routers, if you remove the redundant Switch Interface Board (SIB) after upgrading Junos OS from Release 17.4R1 or Release 17.2X75-D90 to a later release, then an alarm is not generated. This is a known behavior and has no impact on the performance of the router.

SEE ALSO

New and Changed Features | **245**

Changes in Behavior and Syntax | **255**

Known Issues | **264**

Resolved Issues | **269**

Documentation Updates | **276**

Migration, Upgrade, and Downgrade Instructions | **276**

Product Compatibility | **282**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 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.

Forwarding and Sampling

- The **show firewall filter** command does not display policers counters for filters that reference three-color policers. PR1364673

General Routing

- 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-dump. This condition is difficult to reproduce. As a workaround, the interrupt code is optimized to avoid the unnecessary call to prevent the issue. PR1208536
- The following error messages occur during GRES and unified ISSU: syslog errors @ **agentd_rts_async_rtbm_msg : FLM : Failed to create private**. PR1232636
- On the third-generation FPCs of PTX Series routers (PTX3000, PTX5000 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 a CFP2-DCO-T-WDM transceiver installed, after repeated configuration rollbacks, the link sometimes takes a long time to come up. PR1301462

- When a CFP2-DCO-T-WDM-1 transceiver is plugged in a PTX Series PIC, after the FPC restarts, sometimes carrier frequency offset TCA is increased even when TCA not enabled. PR1301471
- Next Generation Routing Engine (NG-RE) with models RE-S-X6-64G, RE-S-2X00x6 and RE-PTX-X8-64G on MX or PTX platform may encounter a transient system freeze of the Linux based host (VMHost) for about 20-35 seconds, causing protocol flaps, FPC restart and mastership switch between Routing Engines. Due to incorrect handling of the disk IO commands, a disk I/O timeout is reported and the system will recover by resetting of the solid-state drives (SSD) channel. The system will continue to operate correct after such an event. PR1312308
- On a PTX Series router with a third-generation FPC, an error message is displayed when the FPC goes online or offline. PR1322491
- On PTX Series devices with TQ-chip cards (for example, FPC1 or FPC2) and class of service (CoS) used, a high-priority queue might not get the entire configured bandwidth. PR1324853
- Protocol-based classification for host-bound traffic does not work. Due to the missing classification, protocol-specific (BGP and OSPF, etc.) statistics are not accounted for correctly. PR1328631
- On a 30-Port MACsec-enabled line card (LC1101-M - 30C / 30Q / 96X) of PTX10008 chassis, when the **exclude-protocol lacp** configuration statement at the **[edit security macsec connectivity-association connectivity-association-name]**, hierarchy level is deleted or deactivated, the LACP Protocol's **Mux State** shown under the output of the CLI command **show lacp interface**, might remain as "attached" or "detached" and might not transition to "distributing" state.PR1331412
- Some default Routing Engine sensors are subscribed as part of default j-insight package. PR1339329
- The same port range (0..19) is used for both PIC 0 and PIC 1. PR1342081
- PTX3000 reports CCL (chip-to-chip link) CRC errors when FPC3-SFF-PTX-1X is taken offline through a CLI command or by pressing the offline button. A syslog error is generated by the FPC just before it goes offline, so there is no detectable traffic loss. PR1348733
- On Next Generation Routing-Engine (NG-RE), a failure of the Hardware Random Number Generator (HWRNG) will leave the system in a state where there are not enough entropy available to operate. PR1349373
- The host path statistics of the Routing Engine and Packet Forwarding Engine might not match. PR1353699
- If output firewall filter is configured with the **syslog** or log option, the host interface might be wedged on PTX1000, PTX5000 and PTX10000. The change in this PR is to add the warning but does not prevent the problem which the host interface stop sending packets. This condition might occur if all below conditions are met:

  1) Packet which is hitting the filter term should be less than 128 bytes

  2) Output firewall filter has syslog, log or port-mirror & accept action. Sample configuration for V4 & V6:

  set interfaces<interface name> unit<unit> family inet filter output <filter-V4>

  set firewall family inet filter <filter-V4> term 1 then log

set firewall family inet filter <filter-V4> term 1 then accept

set interfaces <interface name> unit family inet6 filter output <filter-V6>

set firewall family inet6 filter <filter-V6> term 1 then log

set firewall family inet6 filter <filter-V6> term 1 then acceptPR1354580

- On PTX1000-M20C, while configuring MAC-Sec over a 40g interface, it might lead to invariable blocking of these ports leading to traffic drop. To circumvent this issue, it is advisable to configure the 40g speed over the interfaces and reboot the box before the MacSec configuration. PR1357849
- With aggregated Ethernet flap, OSPF session establishment might take more time than average to converge. PR1359343
- Committing aggregated Ethernet configuration might lead to minor errors. PR1365355
- When the TIC goes offline and comes back online, MPLS bidirectional traffic flow might stop working. PR1367920
- Frequent speed changes on interface ports might cause the relevant port IFD not to be created. PR1367946
- Some harmless log messages are suppressed on the backup SPMB. PR1369731
- You might not be able to stop the ZTP bootstrap when a PTX10016 or PTX10008 router with more number of line cards is powered ON with the factory default configuration. PR1369959
- When a Routing Engine reboots and comes up again, it sends gratuitous ARP packets to the internal interfaces in order to advertise its MAC address. These packets get into the UKERN running on the FPC, which drops these packets. The messages are displayed just before these packets are dropped. These errors are harmless and do not disrupt the working of any feature. PR1374372
- When Jflow sampling is enabled and flows are sampled through aggregate bundles, the following harmless error logs are generated: **[Tue Oct 30 18:17:40.648 LOG: Info] expr_get_local_pfe_child_ifl: cannot find child ifl of agg ifl 74 for this fpc [Tue Oct 30 18:17:40.648 LOG: Info] flowtb_get_cpu_header_fields: Failed to find local child ifl for 74 [Tue Oct 30 18:17:40.648 LOG: Info] fpc0 cannot find stream on [hostname]** .PR1379227
- If multiple LLDP sensors are exported together and part of their keys overlap, the data for these sensors might not get exported. PR1382691
- On PTX Series routers or QFX10002, QFX10008, and QFX10016, an auto correctable non-fatal hardware error on the Provider Edge chip (which is ASIC on PTX1000, PTX10002, and QFX10002, the third-generation FPC on PTX3000 and PTX5000, and the line card on PTX10008, PTX10016, QFX10008, and QFX10016) is reported as 'FATAL' error, and therefore the relevant Packet Forwarding Engine (PFE) is 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
- Deviation and augmentation updated for IS-IS telemetry. Updated in: Junos OS Releases 18.1, 18.2, 18.3, 18.4, 19.1, 19.2, and 19.3 branches. PR1408151

- When a 100g QSFP is inserted into FPC on PTX, all the other interfaces on that FPC and the other FPCs might flap, since these interfaces are configured the smaller **pdu-interval** value of LFM. PR1408204
- When forwarding chain is unilist_1->indirect-next-hop->unilist_2, any change in unilist_2 active member list is absorbed by indirect-next-hop in the chain and the change will not be back propagated to top-level unilist_1. If a link flap it will cause indirect-next-hop pointing to unilist_2 stuck with weight 65535 and further causing traffic blackholing. PR1409632
- On PTX10002 devices, if the **chassis-control** process restarts, Express ASICs are not initialized, leading to packet drops on the output queue. PR1414434
- While committing the configuration, the following error message is displayed: error: **mustd trace init failed**. PR1423229
- On routers and switches with Link Aggregation Control Protocol (LACP) enabled, deactivating a remote aggregated Ethernet member link changes the LACP state of the local member link to Detached state. The detached link is invalidated from the Packet Forwarding Engine's AE-Forwarding Table. If the device is rebooted with this state, all the member links are enabled in PFE AE-Forwarding Table irrespective of the LACP states, resulting in packet drop. PR1423707
- PTX10000/LC1101: when an interface is configured with jumbo frames support (for example, MTU = 9216), the effective MTU size for locally sourced egress traffic is 24 bytes less than the expected value. This issue is confined to locally originated traffic only and does not affect transit traffic. PR1428094
- Configuring IP Flow Information Export (IPFIX) on a device with LC1104 or LC1105 line cards might trigger major/Fatal ASIC errors, and the Packet Forwarding Engines might shut down. PR1429419
- The timestamp reported for packet arrival in NetFlow records reports inaccurate time due to a synchronization issue with NTP. PR1431498

Infrastructure

- A file system corruption might create a kernel core file. The Routing Engine reboots with the message; **ffs_blkfree: freeing free block**. PR1028972
- Junos packages might be incorrectly registered as **unsupported**. PR1427344

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

MPLS

- When the rpd daemon terminates, the process of signaling the deletion of all RSVP LSPs might take so long that a watchdog timer is triggered, resulting in the generation of an rpd core file. PR1257367
- Due to an error with the optimization timer, a particular check fails when the exponential increase function is called. PR1416948

Platform and Infrastructure

- Use groups re0/re1 to configure the Routing Engine-specific management interface. PR1375012

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 the source address, BFD packets are sent with the source address of the outgoing interface instead of the loopback address. Due to this issue, the micro BFD session might not come up. PR1370463

SEE ALSO

New and Changed Features | **245**

Changes in Behavior and Syntax | **255**

Known Behavior | **261**

Resolved Issues | **269**

Documentation Updates | **276**

Migration, Upgrade, and Downgrade Instructions | **276**

Product Compatibility | **282**

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

General Routing

- The PTX5000 router experiences more than 50 percent multicast traffic packet drop. PR1339481
- MPLS LSP statistics are not shown in the output of the **show mpls lsp ingress statistics** command. PR1344039
- The traffic loss duration during FRR link-protection is between 25 ms and 150 ms. PR1355953
- **slu.l2_domain_lookup_failure** traps might be observed when using sampling on FPC-P1/FPC-P2. PR1368381
- Unexpected incrementing of counters on the interface. PR1370062
- Power usage ST components in PTX5000 do not work as intended. PR1372369
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent major alarm. PR1384435
- Packet drop might be seen in lower priority queues on PTX Series or QFX10000 Series platforms. PR1385454
- lcmd core files are generated and FPC is restarted. PR1391443
- Agentd sensor transmits multiple interface telemetry statistics per FPC slot. PR1392880
- The **show chassis fpc** command on PTX Series and QFX10000 Series routers shows incorrect buffer memory utilization. PR1397612
- 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 can be disabled on an FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- The TCP connection between ppmd and ppman might be dropped due to a kernel issue. PR1401507
- The log message **JAM HW data base open failed for ptx5kpic_3x400ge-cfp8** is generated 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 more than 10 minutes after a reboot of the chassis. PR1404611
- On PTX3000 and PTX5000, backup CB's chassis environment status keeps displaying after the backup CB is removed or power to PEM is lost. PR1405181
- 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
- Layer 2 VPN might flap repeatedly after the link up between PE and CE devices. PR1407345
- The port at FPC (for example, JNP10K-LC1101) might fail to come up. PR1409585
- Hostname does not update at FPC shell after system configuration change on CLI. PR1412318
- Junos PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. PR1412659
- The L2circuit egress Provider Edge 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
- RX alarms are not set as according to the threshold value configured for the DCO Tunable Optics. PR1419204
- Error messages might be seen on PTX10000/QFX10000 platforms during DFE tuning. PR1421075
- Virtual Chassis might become unstable and FXPC might generate core files when there are a lot of configured filter entries. PR1422132
- Packet Forwarding Engine wedge might be observed after performing the **show forwarding-options load-balance** command.PR1422464
- 4x10G interfaces on PTX3000/PTX5000 FPC type 3 might not come up after frequently flapping for a large amount of time. PR1422535
- Specific interface on P3-15-U-QSFP28 PIC card remains down until another interface comes up. PR1427733

Infrastructure

- The **request system recover oam-volume** command might fail on PTX Series routers. PR1425003

Interfaces and Chassis

- The syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** is seen upon LFM related configuration commit on aggregated Ethernet interfaces. PR1423586
- Some ports on PTX Series routers might remain down after rebooting the FPC/device at remote side. PR1429315

MPLS

- An RSVP-signaled LSP might stay in down state after a link in the path. flaps. PR1384929
- The rpd might crash when an 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
- LDP routes might flap if committing any configuration changes. PR11416032
- Bypass dynamic rsvp lsp tears down too soon when being used for protecting ldp lsp with knob **dynamic-rsvp-lsp**. PR1425824

Platform and Infrastructure

- Some files are missing during log archiving. PR1405903

Routing Protocols

- RPD core files are generated on backup Routing Engine during neighborship flap when using **authentication-key** with size larger than 20 characters. PR1394082
- Syslog message is seen whenever prefix-sid coincides with the node-sid. PR1403729
- The rpd memory leak might be seen in IS-IS segment routing scenario. PR1404134
- Dynamic routing protocol flapping with vmhost Routing Engine switchover is seen on new generation Routing Engine. PR1415077

VPNs

- In a specific CE device environment in which asynchronous notification is used, after the link between the PE and CE devices goes up, the Layer 2 circuit flaps repeatedly. PR1282875

Resolved Issues: 18.2R2

General Routing

- On a PTX1000, upgrade from Junos OS Release 16.1X65D45 to Junos OS Release 17.3-20170721 fails frequently with sampling enabled. PR1296533
- Status LED on the chassis does not light up on QFX10002-60C. PR1332991
- Tc_count counters for a filter with the **scale-optimized** statement enabled do not increment. PR1334580
- Members of IPv4 unilist next hops might be stuck in "Replaced" state after interface flapping. PR1336201
- On QFX10000 platforms, NETCONF SSH TCP port 830 traffic hitting host path/unclassified queue. PR1345744
- FPC reboots a few minutes after the configuration is loaded. PR1346467
- Packet might be dropped by RPF during Routing Engine switchovers. PR1354285
- Unable to commit Junos OS configuration during the ZTP process, and the ZTP process stop completed. PR1358919
- Multicast replication traffic might be lost on an aggregated Ethernet bundle interface after one member link goes down. PR1359974
- The route might be stuck after BGP neighbor and route flapping. PR1362560
- Traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with “Link-Layer-Down" flag set. PR1365263
- PTX IPLC might not boot up with multiple J-UKERN crashes. PR1365791
- The 'Normal discards' Packet Forwarding Engine statistics traffic counter might increase at a higher rate when Inline J-flow or S-Flow is enabled. PR1368208
- JNP hardware running Junos for QFX software messages is continuously getting flooded with **dcpfe_pd[4235]: et-0/1/5:2: Signal lost. Macsec rx 0**. PR1368969
- The commit or commit check operation might fail because of the error **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- Packets might be dropped after a filter is deleted from an interface. PR1372957
- Inline BFD keep flapping when inline samping is configured. PR1376509
- Traffic might be dropped on third-generation FPCs on PTX. PR1378392
- Layer 3 VPN traffic might be dropped because one core-facing interface is down. PR1380783

- BFD sessions bounced on FPCs that are not taken offline. PR1383703
- CPSM daemon memory leak in VM host. PR1387903
- Forwarding issue on mixed link-speed aggregated Ethernet interface after FPC reloads. PR1390417

Infrastructure

- The FPC might go down on some VM-host-based PTX Series or QFX Series devices. PR1367477

Interfaces and Chassis

- Major error **PE Chip:pe0[0]: IPW: oversize_drop error** seen on FPC. PR1375030

MPLS

- LSP with **auto-bandwidth** enabled goes down during HMC error condition. PR1374102

Platform and Infrastructure

- Junos OS: Next hop index allocation failed: private index space exhausted because of incoming ARP requests to the management interface (CVE-2018-0063). PR1360039

Routing Protocols

- Rpd core files might be generated during telemetry streaming. PR1347431

Resolved Issues: 18.2R1

General Routing

- Remove **show chassis spmb** command and response. PR1244059
- For MTRE devices using telemetry, **restart na-grpc-server** and **restart na-mqtt** do not work. PR1284121
- For BGP-LU multipath routes, if there is a forwarding-table export policy configured to reject such routes, then rpd might might crash during next-hop installation. PR1297044
- Interfaces might go down when the Packet Forwarding Engine encounters **TOE::FATAL ERROR**. PR1300716
- The FPC is being reported as down in chassisd logs related to streaming telemetry, eventhough the FPC is online. PR1300795
- A third-generation FPC (FPC3-SFF-PTX) might not boot on a PTX3000 with the Control Board or Routing Engine. PR1303295
- Internal latency is high during initial subscription of sensors. PR1303393
- The mgd might process crash when the Ephemeral database is used. PR1305424
- Packet Forwarding Engine error messages are flooding as **expr_sensor_update_cntr_to_sid_tree** after delete and rollback of **protocols isis source-packet-routing node-segment**. PR1309288
- Need to suppress chassis alarm for switched off PEM. PR1311574

- The SIB LED on the front panel display is green and remains steadily lit even before an SIB comes online. PR1311632
- When the user changes the PIC or port speed, an alarm is raised and user intervention is required. PR1311875
- Memory leak in the chassisd process occurs while streaming telemetry subscriptions are active. PR1315672
- Packet Forwarding Engine packet drop is seen on the PTX5000 when there is a 100-ms RTT delay between the DUT and the collector. PR1316429
- On the PTX10000, for 100G LR4 Optics with part number 740-061409, need to change **show chassis hardware** display to QSFP-100G-LR4-T2. PR1322082
- The rpd might crash when an OpenConfig package is upgraded with JTI streaming data in the background. PR1322553
- On PTX1000, MX204, MX10003, or QFX10002-60C, the local time on the FPC might be different from the local time on the Junos VM or VM host. PR1325048
- The GRE traffic is not de-encapsulated by the firewall filter. PR1325104
- Firewall filter is not supported on aggregated Ethernet. PR1325237
- PTX Series MKA sessions are not coming up after changing CA parameters such as - **transmit-interval**, and **key-server-priority**. PR1325392
- MPLS traceroute fails across PTX Series platform. PR1327609
- Unsupported features need to be removed or disabled under CLI **set vlans <vlan_name>**. PR1328219
- Unsupported options need to be disabled under CLI **set interfaces <interface_name> unit 0 family ethernet-switching interface-mode trunk**. PR1328507
- Link instability occurs after link-down event on PTX Series device. PR1330708
- Traffic stops flowing out of ae70 after some FPC restart iterations. PR1335118
- PTX5000 FPC might reboot in certain rare scenarios when interface-specific policer is configured. PR1335161
- Disabling a breakout 10G port on et-0/0/5 will unexpectedly disable another breakout 10G port on et-0/0/5. PR1337975
- FPC/FPC2/FPC E on PTX Series device does not forward traffic. PR1339524
- Link goes down on PTX3000/PTX5000 with FPC3 inserted after router reboot or link flap. PR1340612
- On the PTX1008, the 30-Port Coherent Line Card (DWDM-IC) does not come up. PR1344732
- No DHCP service or configuration is running after the system has cleared. PR1347730
- Sensors are not getting cleared up after doing Routing Engine switchover. PR1347779

- Threshold is not getting configured correctly in PTX Series device when threshold is configured using scope and category options. PR1350841
- BFD sessions do not come up on PTX3000. PR1352112
- Flabels might get exhausted after multiple Routing Engine switch-over. PR1354002
- The interface of 15 100G ports PIC might delay 60 seconds to come up. PR1357410

Infrastructure

- The ixlv interface statistics are not accounted for properly. PR1313364

Interfaces and Chassis

- On the PTX3000, failed to check CFM neighbors wrt **show oam ethernet connectivity-fault-management interfaces ae0.0 extensive**. PR1335305
- The transportd process might crash when an SNMP query is performed on jnxoptIfOChSinkCurrentExtTable with an unsupported interface index. PR1335438

MPLS

- Traffic drop is seen during NSR switchover for RSVP P2MP provider tunnels used by MVPN. PR1293014
- Traffic loss occurs for static LSP configured with the **stitch** command. PR1307938
- The rpd might crash on the backup Routing Engine due to memory exhaustion. PR1328974
- MPLS LSP statistics are not shown in cli command **show mpls lsp ingress statistics**. PR1344039

Platform and Infrastructure

- DCD Microbfd seems to be failing in dcd_commit_check log file even when BFD is not configured. PR1300796
- Traffic might be silently dropped and the following message might be seen: **JPRDS_NH:jprds_nh_alloc(),651: JNH[0] failed to grab new region for NH messages**. PR1357707
- Unable to commit junos configuration during the ZTP process and ZTP process stop completed. PR1358919

Routing Protocols

- The rpd might constantly consume high CPU resources in a BGP setup. PR1315066
- The primary path of MPLS LSP might switch to another address. PR1316861
- The rpd process might crash continuously on both Routing Engines when **backup-spf-options remote-backup-calculation** is configured in IS-IS protocol. PR1326899
- Protocol churn will create rpd crash. PR1341466

SEE ALSO

New and Changed Features | 245
Changes in Behavior and Syntax | 255
Known Behavior | 261
Known Issues | 264
Documentation Updates | 276
Migration, Upgrade, and Downgrade Instructions | 276
Product Compatibility | 282

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 documentation for the PTX Series.

SEE ALSO

New and Changed Features | 245
Changes in Behavior and Syntax | 255
Known Behavior | 261
Known Issues | 264
Resolved Issues | 269
Migration, Upgrade, and Downgrade Instructions | 276
Product Compatibility | 282

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

- Upgrade and Downgrade Support Policy for Junos OS Releases | 277
- Upgrading a Router with Redundant Routing Engines | 277
- Basic Procedure for Upgrading to Release 18.2 | 278
- Installing the Software on PTX10002-60C Routers | 281

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.

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 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 acting as the backup Routing Engine.

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

Basic Procedure for Upgrading to Release 18.2

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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

To download and install Junos OS Release 18.2R3:

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:** After you install a Junos OS Release 18.2R3 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.

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 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.

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

```
user@host> request system software add validate reboot
source/jinstall-18.2R3.SPIN-domestic-signed.tgz
```

All other customers, use the following command:

```
user@host> request system software add validate reboot
source/jinstall-18.2R3.SPIN-export-signed.tgz
```

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

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

- **scp://*hostname*/*pathname*** (available only for Canada and U.S. version)

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

Adding the **reboot** command reboots the router after the upgrade is validated and installed. When the reboot is complete, the router displays the login prompt. The loading process 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 18.2 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

Installing the Software on PTX10002-60C Routers

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 PTX10002-60C 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-ptx-x86-64-18.2R3.9.tgz
```

If the installation 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-ptx-x86-64-18.2R3.9.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
```

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for the QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.2R3 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 Junos OS main release and the maintenance releases for QFX Series.

> **NOTE:** The following QFX Series platforms are supported in Release 18.2R3: QFX5100, QFX5110, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016.

Release 18.2R3 New and Changed Features

- There are no new features or enhancements to existing features for QFX Series in Junos OS Release 18.2R3.

Release 18.2R2 New and Changed Features

- There are no new features or enhancements to existing features for QFX Series in Junos OS Release 18.2R2.

Release 18.2R1 New and Changed Features

Hardware

- **QFX10000-30C-M line card supports channelization (QFX10008 and QFX10016 switches)**—Starting in Junos OS Release 18.2R1. 40-Gigabit Ethernet ports on the QFX10000-30C-M line card can be channelized to 10-Gigabit Ethernet. When ports are in channelization mode, every fifth port is disabled. [See QFX10000-30C-M Line Card.]
- **Support for JNP-QSFP-100G-BXSR transceiver (QFX5200)**—Starting in Junos OS Release 18.2R1, the QFX5200 switches support the JNP-QSFP-100G-BXSR transceiver. The 100 Gigabit bidirectional transceiver has a dual transmitter/receiver that allows it to transmit and receive data through a single optical fiber. Each bidirectional transceiver has two LC receptacles that receive and transmit on different optical wavelengths. The wavelength of the input optical signal needs to match the receive wavelength of the pairing transceiver. For example, if transceiver A has a transmit wavelength of 850 nm and a receive wavelength of 900 nm, then the pairing transceiver B should have a matching receive wavelength of 850 nm and a transmit wavelength of 900 nm.
  [See the Hardware Compatibility Tool.]

Authentication Access Control

- **Enhancement to NTP authentication method (QFX5110, QFX10000)**— Starting in Junos OS Release 18.2R1, Junos OS supports NTP authentication for both SHA-1 and SHA2-256, in addition to the existing NTP authentication method, MD5. You can now choose from among MD5, SHA-1, and SHA2-256 for synchronizing the clocks of Juniper Network routers, switches, and other security devices on the Internet.

Using SHA-1 instead of MD5 improves the security of devices with very little impact to timing, while using SHA2-256 provides an increase in security over SHA-1.

> **NOTE:** By default, network time synchronization is unauthenticated.

To implement authentication, use **set authentication-key *key_number* type** at the **[edit system ntp]** hierarchy level.

- To enable SHA-1 authentication, use **set authentication key *key_number* type sha1 value *password*** at the **[edit system ntp]** hierarchy level.
- To enable SHA2-256 authentication, use **set authentication key *key_number* type sha256 value *password*** at the **[edit system ntp]** hierarchy level.

[See authentication-key and Configuring NTP Authentication Keys.]

EVPN

- **Support for firewall filtering and policing on EVPN-VXLAN traffic (QFX5100 and QFX5110)**—Starting with Junos OS Release 18.2R1, you can configure firewall filters and policers on VXLAN traffic in an EVPN topology. Firewall filters provide rules that define whether to accept or discard packets that are transiting an interface. Policing, or rate limiting, lets you control the amount of traffic that enters the switch and further determines the actions to be taken when the traffic exceeds the defined limit. You configure firewall filters at the **[edit firewall]** hierarchy level. For each firewall filter that you apply to a VXLAN, you can specify **family ethernet-switching** to filter Layer 2 (Ethernet) packets, or **family inet** to filter on IRB interfaces. The IRB interface acts as a Layer 3 routing interface to connect the VXLANs in one-layer or two-layer IP fabric topologies. You can only apply firewall filters and policers only on CE-facing interfaces in the ingress direction (traffic entering the VXLAN). For IRB interfaces, you can apply filtering only at the ingress point of non-encapsulated frames routed through the IRB interface.

  This feature is not supported on a QFX5100 Virtual Chassis in an EVPN-VXLAN topology.

  [See Understanding VXLANs and Overview of Firewall Filters.].

- **IPv6 data traffic support through an EVPN-VXLAN overlay network (QFX5110 switches)**—Starting with Junos OS Release 18.2R1, 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 Layer 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. To enable IPv6 data traffic support, you configure the IRB interfaces on all Layer 3 VXLAN gateways with the same IPv4 and IPv6 anycast virtual gateway addresses (VGAs). To support this feature, no other IPv6 configuration is required in the underlay or overlay networks.

  (The feature described above is documented but not supported on QFX5110 switches in Junos OS Release 18.2R1.)

[See Routing IPv6 Data Traffic Through an EVPN-VXLAN Network with an IPv4 Underlay.]

- **Support for OSPF, IS-IS, BGP, and static routing on IRB interfaces in EVPN-VXLAN networks (QFX5110)**—Starting in Junos OS Release 18.2R1, you can configure OSPF, IS-IS, BGP, and static routing with bidirectional forwarding detection (BFD) on an IRB interface that is used as a routed interface in EVPN. This configuration allows protocol adjacencies to be established between an IRB interface on a Layer 3 gateway and a CE device and between an IRB interface on a Layer 3 gateway and a CE device connected to a Layer 2 leaf device in an EVPN-VXLAN network.

  [See Supported Protocols on an IRB Interface in EVPN-VXLAN .]

- **Support for IS-IS on IRB interfaces in EVPN-VXLAN networks (QFX 10000)**—Starting in Junos OS Release 18.2R1, you can configure IS-IS on an IRB interface that is used as a routed interface in EVPN. This configuration allows protocol adjacencies to be established between an IRB interface on a Llayer 3 gateway and a CE device and between an IRB interface on a Layer 3 gateway and a CE device connected to a Layer 2 leaf device in an EVPN-VXLAN network.

  [See Supported Protocols on an IRB Interface in EVPN-VXLAN .]

- **NOTE:** QFX5110 and QFX5200 switches do not currently support the pop functionality, which has the following implications for this feature:

  - The following use cases are not supported:
    - Traffic Pattern 1: Popping an S-VLAN tag
    - Traffic Pattern 4: Popping and later pushing an S-VLAN tag
  - Without the pop functionality, this feature does not actually support the tunneling of Q-in-Q traffic through an EVPN-VXLAN overlay network. The functionality that is currently supported is flexible VLAN tagging.

Tunneling Q-in-Q traffic through an EVPN-VXLAN overlay network (QFX5110 and QFX5200 switches)—Starting with Junos OS Release 18.2R1, QFX5110 and QFX5200 switches that function as Layer 2 VXLAN tunnel endpoints (VTEPs) can tunnel single-tagged and double-tagged Q-in-Q packets through an Ethernet VPN-Virtual Extensible LAN (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 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 QFX5110 and QFX5200 switches do not support the pop and push actions with a configured range of VLANs.

The ingress and egress VTEPs support the tunneling of Q-in-Q packets and the Q-in-Q actions in the context of specific traffic patterns.

[See Examples: Tunneling Q-in-Q Traffic in an EVPN-VXLAN Overlay Network.]

Interfaces and Chassis

- **Channelization support on QFX10000-30C-M line cards (QFX10008 and QFX10016 switches)**–Starting in Junos OS Release 18.2R1, you can channelize the 40-Gbps port speeds of the QFX10000-30C-M line card into four independent data channels of 10-Gbps. The Media Access Control Security (MACsec) ports auto sense the transceiver and set the port to the default (non-channelized) mode D. By changing a port to mode A (channelized), the associated Packet Forwarding Engine reboots the four ports that it controls and disables one port.

Junos Telemetry Interface

- **Packet Forwarding Engine sensors for the Junos Telemetry Interface (QFX5100, QFX5110, and QFX5200 Switches)** —Starting with Junos OS Release 18.2R1, you can export Packet Forwarding Engine statistics through the Junos Telemetry Interface using native sensors. Native sensors export data close to the source, such as the line card or network processing unit (NPU), using the User Datagram Protocol (UDP).

  The native sensors listed in Table 1 are supported.

  Table 2: Supported Packet Forwarding Sensors

Sensor	Exports
/junos/system/linecard/qmon-sw/ **TIP:** This sensor is only available on QFX5000 Series Switches.	Statistics for congestion and latency monitoring
/junos/system/linecard/interface/logical/usage/	Logical interface statistics
/junos/system/linecard/firewall/	Filter statistics
/junos/system/linecard/interface/	Physical interface statistics
/junos/services/label-switched-path/usage/	Label-switched paths (LSP) statistics
/junos/system/linecard/cpu/memory/	Network Processing Unit (NPU)/Line Card memory

  For streaming statistics through UDP, all parameters are configured at the **[edit services analytics]** hierarchy level.

  Support for the Junos Telemetry Interface was introduced on QFX10000 and QFX5200 switches in Junos OS Release 17.2R1.

[See sensor and Configuring a Junos Telemetry Interface Sensor (CLI Procedure).]

- **Streaming OpenConfig data from Routing Engine sensors over UDP in protobuf format (QFX Series)**–Starting in Junos OS Release 18.2R1, you can stream OpenConfig-based sensor data from Routing Engine sensors by using the Junos Telemetry Interface (JTI). JTI enables you to stream the OpenConfig sensor data in gRPC/protobuf format rather than in key/value pairs. Using this format is more efficient and makes the messages smaller.

  [See Overview of the Junos Telemetry Interface.]

Port Security

- **IPv6 Router Advertisement guard (RA guard) (QFX5100/QFX5110/QFX5200)**—Starting with Junos OS Release 18.2R1, IPv6 RA guard is supported on QFX5100, QFX5110, and QFX5200 switches. router advertisement guard protects networks 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 the messages based on whether they meet certain criteria, which is configured on the switch as a policy. RA guard inspects the router advertisement 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.]

- **Client link-layer address option 79 for DHCPv6 (QFX5100/QFX5100-VC, QFX5110/QFX5110-VC, QFX5200, QFX10002, QFX10008, QFX10016)**—Starting in Junos OS Release 18.2R1, you can configure DHCPv6 option 79 to insert the DHCPv6 client link-layer address in the header of the DHCPv6 RELAY-FORWARD message that is sent from the client to the upstream device. The client link layer address can be used along with other identifiers to associate DHCPv4 and DHCPv6 messages from a dual-stack client.

  [See Inserting the DHCPv6 Client MAC Address Option (Option 79) In DHCPv6 Packets].

Restoration Procedures Failure

- **Device recovery mode introduced in Junos OS with upgraded FreeBSD (QFX Series)**—Starting in Junos OS Release 18.2R1, for devices running Junos OS with upgraded FreeBSD, provided you have saved a rescue configuration on the device, an automatic device recovery mode is triggered if the system goes into amnesiac mode. In this new process, the system automatically retries to boot with the saved rescue configuration. The system displays the banner **Device is in recovery mode** in the CLI (in both the operational and configuration modes).In earlier releases of Junos OS, there is no automatic process to recover from amnesiac mode; therefore a user with load and commit permission must log in using the console and fix the issue in the configuration before the system can reboot.

  [See Saving a Rescue Configuration File.]

Routing Protocols

- **Remote LFA support for LDP in IS-IS and OSPF (QFX5100, QFX5110, QFX5200)**—Beginning with Junos OS Release 18.2R1, you can configure a remote loop-free alternate (LFA) to extend the backup provided by the LFA in an IS-IS or OSPF network. This feature is useful especially for Layer 1 metro rings where the remote LFA is not directly connected to the point of local repair (PLR). The existing LDP implemented for the MPLS tunnel setup can be reused for the protection of IS-IS and OSPF networks and subsequent LDP destinations, thereby eliminating the need for RSVP-TE backup tunnels for backup coverage.

  To configure remote LFA over LDP tunnels in an IS-IS network, include the **remote-backup-calculation** statement at the **[edit protocols isis backup-spf-options]** hierarchy level and the **auto-targeted-session** statement at the **[edit protocols ldp]** hierarchy level.

  [See Example: Configuring Remote LFA over LDP Tunnels in IS-IS Networks. and Example: Configuring Remote LFA Over LDP Tunnels in OSPF Networks.]

Security

- **Support for CCC firewall filters (QFX10000 switches)**—Starting with Junos OS Release 18.2R1, you can configure inbound and outbound firewall filters with counter and policer actions on Layer Layer 2 circuit cross-connect (CCC) traffic (**family ccc**). This feature is beneficial if you use Layer 2 point-to-point circuits to connect customers between sites and want to use policers to apply limits to traffic flowing over CCC circuits. You configure Layer 2 firewall filters at the **[edit firewall filter family ccc]** hierarchy level.

  [See CCC Overview and Firewall Filter Match Conditions for Layer 2 CCC Traffic.]

Software Installation and Upgrade

- **Zero Touch Provisioning (QFX10008 and QFX10016 switches)**—Starting with Junos OS Release 18.2R1, you can use Zero Touch Provisioning to provision new Juniper Networks switches in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, it attempts to automatically upgrade the Junos OS software and install a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network

  [See Zero Touch Provisioning.]

System Management

- **Support for the Precision Time Protocol (PTP) AES67, SMPTE ST-2059-2, and AES67+SMPTE profiles (QFX5110-48S and QFX5200 switches)**—Starting in Junos OS Release 18.2R1, you can enable the AES67, SMPTE ST-2059-2, and AES67+SMPTE profiles to support video applications for capture (for example, cameras), video edit, and playback to be used in professional broadcast environments. The PTP standard allows multiple video sources to stay in synchronization across various equipment by providing time and frequency synchronization to all devices. These profile support PTP over IPv4 multicast and ordinary and boundary clocks.

  To configure the AES67, SMPTE ST-2059-2, and AES67+SMPTE profiles, enable one of the **aes67**, **smpte**, or **aes67-smpte** statements at the **[edit protocols ptp profile-type]** Junos OS CLI hierarchy.

  See [Understanding the PTP Media Profiles.]

- **Zero Touch Provisioning (QFX10002-60C switches)**—Starting with Junos OS Release 18.2, Zero Touch Provisioning allows you to provision new Juniper Networks routers in your network automatically without manual intervention. When you physically connect a switch to the network and boot it with a default configuration, the switch attempts to automatically upgrade the Junos OS software image and install a configuration file from the network. The switch uses information that you configure on a Dynamic Host Configuration Protocol (DHCP) server to locate the necessary software image and configuration files on the network. If you do not configure the DHCP server to provide this information, the switch boots with the preinstalled software and default configuration. The Zero Touch Provisioning process either upgrades or downgrades the Junos OS version.

  [See Understanding Zero Touch Provisioning.]

- **New tool to detect high CPU utilization (QFX Series)**—Starting in Junos OS Release 18.2R1, a flight recorder tool is introduced to gather historical data on when the CPU utilization on a device was high and what processes caused the high utilization. The tool collects snapshots of data, enabling detection of high CPU usage and faster resolution of issues.

  Because some of the high CPU utilization cases are intentional or expected, you can enable and disable the flight recorder tool to avoid false alarms.

  [See request flight-recorder set high-cpu and show flight-recorder status.]

VLAN Infrastructure

- **Flexible Ethernet support (QFX10000 Switches)**—Starting in Junos OS Release 18.2R1, you can configure inet, inet6, or VLAN circuit cross connect (CCC) connections on a physical or aggregated Ethernet interface. This configuration enables you to set different forwarding rules for tagged and untagged traffic on the same interface. For example, you can forward tagged packets over the Layer 2 circuit and route untagged traffic normally in the native VLAN mode.

  All logical devices that are under the flexible VLAN tagging are identified by their VLAN ID configuration. For untagged traffic, the association to the corresponding logical device is derived using the native VLAN ID configuration on the physical device. For traffic without a VLAN tag, the default VLAN ID (or native VLAN ID) is used to derive the Layer 2 domain.

SEE ALSO

Changes in Behavior and Syntax | 291

Known Behavior | 297

Known Issues | 300

Resolved Issues | 308

Documentation Updates | 326

Migration, Upgrade, and Downgrade Instructions | 326

Product Compatibility | 340

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 18.2R3 for the QFX Series.

EVPN

- **New options in show evpn instance command (QFX series)**—Starting in Junos OS Release 18.2R3, you can use the **show evpn instance esi-info** command to display only the ESI information for a routing instance and **show evpn instance neighbor-info** to display only the IP address of the EVPN neighbor for a routing instance. Information associated with the ESI, such as the route distinguisher, bridge domain, and IRB are filtered out.

High Availability (HA) and Resiliency

- **commit fast-synchronize** option not supported for products with single Routing Engine (QFX Series)—Starting in Junos OS Release 18.2R1, Junos OS does not support the configuration option **commit fast-synchronize** at the **[edit system]** hierarchy level for all the products with single Routing Engine for which **chassis redundancy graceful-switchover** is not supported. This option is disabled from the CLI.

Interfaces and Chassis

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (QFX Series)**—In Junos OS Release 18.2R3, 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.
- **The resilient-hash statement is no longer available under aggregated-ether-options (QFX5200 and QFX5210 switches)**—Starting in Junos OS Release 18.2R3, the **resilient-hash** statement is no longer available in the **[edit interfaces *aex* aggregated-ether-options**] hierarchy level. Resilient hashing is not supported on LAGs on QFX5200 and QFX5210.

  [See aggregated-ether-options.]

- **Logical interfaces created along with physical interfaces by default (QFX10000 and QFX5000 switches)**—On the QFX10000 line of switches, logical interfaces are created along with the physical et-, sxe-, xe-, and channelized xe- interfaces. In earlier releases, only physical interfaces are created.

  On the QFX5000 line of switches, by default, logical interfaces are created on channelized xe- interfaces. In earlier releases, logical interfaces are not created by default on channelized xe- interfaces (xe-0/0/0:1, xe-0/0/0:2, and so on), but they are created on et-, sxe-, and nonchannelized xe- interfaces.

- **Commit error when GRE interface and tunnel source interface configured in different routing instances (QFX Series)**—In Junos OS Release 18.2R3, QFX Series switches do not support the configuration of the 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

- **Junos XML protocol <open-configuration> operation no longer emits an uncommitted changes warning (QFX Series)**—Starting in Junos OS Release 18.2R1, the Junos XML protocol **<open-configuration>** operation does not emit an **uncommitted changes will be discarded on exit** warning message when opening a private copy of the candidate configuration. However, Junos OS still discards the uncommitted changes upon closing the private copy.

- **MD5 and SHA-1 hashing algorithms are no longer supported for script checksums (QFX Series)**—Starting in Junos OS Release 18.2R2, 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.

Junos Telemetry Interface

- **Change to the configuration location for gRPC-based sensor subscriptions from an external collector (QFX Series)**—Starting in Junos OS Release 18.2R1, when an external streaming server, or collector, provisions sensors to export data through gRPC on devices running Junos OS, the sensor configuration is committed to the **junos-analytics** instance of the ephemeral configuration database, and the configuration can be viewed by using the **show ephemeral-configuration instance junos-analytics** operational command. In earlier releases, the sensor configuration is committed to the default instance of the ephemeral configuration database.

Layer 2 Features

- **input-native-vlan-push (EX2300, EX3400, EX4600, EX4650, and the QFX5000 line of switches)**—From Junos OS Release 18.2R3, the configuration statement **input-native-vlan-push** at the **[edit interfaces *interface-name*]** hierarchy level is introduced. You can use this statement in a Q-in-Q tunneling configuration to enable or disable whether the switch inserts a native VLAN identifier in untagged frames received on the C-VLAN interface, when the configuration statement **input-vlan-map** with a **push** operation is configured.

  [See input-native-vlan-push.]

MPLS

- When the **no-propagate-ttl** statement is configured on a QFX5200 switch in an MPLS network, the TTL value is not is not copied and decremented on the transit devices during a swap operation. When the switch acts as an ingress device for an LSP, it pushes an MPLS header with a TTL value of 255, regardless of the IP packet TTL. When the switch acts as the penultimate provider switch, it pops the MPLS header without writing the MPLS TTL into the IP packet. PR1368417

Network Management and Monitoring

- **New context-oid option for trap-options configuration statement to distinguish the traps that come from a nondefault routing instance and nondefault logical system (QFX Series)**—In Junos OS Release 18.2R1, a new option, **context-oid**, for the **trap-options** statement enables 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.]

- **Junos OS does not support management of YANG packages in configuration mode (QFX Series)**—Starting in Junos OS Release 18.2R2, 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 18.2R2, 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 18.2R1, 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 the **gtp-tunnel-endpoint-identifier** statement is sufficient for enabling GTP hashing. After enabling, if GTP hashing does not work, we recommend that you capture the packets by using relevant tools and identify the offset value. According to standards, 0x32 is the default header offset value. But, due to some special patterns in the header, the offset value might vary, to say, 0x30, 0x28, and so on. In this cases, use **gtp-header-offset** statement to set a proper offset value. After the header offset value is resolved, run the **gtp-tunnel-endpoint-identifier** command to enable GTP hashing successfully.

  [See gtp-tunnel-endpoint-identifier and gtp-header-offset.]

Security

- **Syslog or log action on firewall lead to packet drops (QFX5000 switches)**—Starting in Junos OS Release 18.2R3, 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 Junos OS Release 18.2R3, a warning message is displayed whenever a firewall term includes the log or syslog option with the accept filter action.

Software Installation and Upgrade

- **New DHCP option introduced for ZTP retry (QFX Series)**—Starting in Junos OS Release 18.2R1, a new DHCP option is introduced to set the timeout value for the file downloads over FTP. If the **transfer-mode** option is set as FTP, the default value for the time out is automatically set as 120 minutes. That is, if the FTP session gets interrupted due to loss of connectivity in the middle of a file transfer, it will timeout

after 120 minutes and ZTP will retry to fetch the file. This value can be overridden using the DHCP option as follows:

```
option NEW_OP.ftp-timeout code 7 = text;
option NEW_OP.ftp-timeout “val”;
```

where **“val”** is the user configurable timeout value in seconds and must be provided within double quotation marks.

Virtual Chassis

- **New configuration option to disable automatic Virtual Chassis port conversion (QFX5100 Virtual Chassis)**—Starting in Junos OS Release 18.2R2, 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 | **283**

Known Behavior | **297**

Known Issues | **300**

Resolved Issues | **308**

Documentation Updates | **326**

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 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.

EVPN

- When a VLAN uses an IRB interface as the routing interface, the VLAN-ID parameter must be set as none to ensure proper traffic routing. This issue is platform-independent. PR1287557
- EVPN-VXLAN implementations support up to 100 EVPN VLAN-based routing instances. If you have more than 100 instances, MAC learning might behave incorrectly. PR1287644

General Routing

- Layer 3 multicast traffic does not converge to 100 percent and a few continuous drops are observed after bringing an interface down and back up again or while an FPC comes online after FPC restart. This behavior is seen when scaling beyond 2000 VLANs or 2000 IRBs with VLAN replication configured. PR1161485
- VLAN tag is removed for inter-VNI traffic on a Layer 3 gateway when the encapsulation or de-encapsulation VLAN configuration or statement is enabled. PR1185295

- When per-packet load balancing is removed or deleted, the next-hop index might change. PR1198092
- Single-bit and multiple-bit ECC errors are not logged on QFX5110 switches. PR1251917
- On the QFX10000-12C-DWDM coherent line card, links might flap when MACsec is enabled on Ethernet interfaces. PR1253703
- On the QFX10000 line of switches, at initialization, the port group module comes up after some time and negative ACKs are seen until the port group module is up. After the port group module is up, negative ACKs are no longer observed. This is an expected behavior due to an aggressive link scan feature introduced in Junos OS Release 17.2. PR1271579
- On the QFX10000 line of switches, with a high scale of 4000 VNIs or 200,000 MACs, or both, if a large configuration change happens with traffic flowing, then forwarding descriptor memory corruption might occur, leading to complete traffic loss on certain ports. The qualification shows that a system with 400 VNIs is stable. However, other configurations such as global MAC count and underlying MPLS LSPs can increase system load. PR1296089
- Traffic drop occurs because of CRC errors when traffic is sent over et-interfaces. PR1313977
- Port LEDs on the QFX5100 do not work. If a device connects to a port on the QFX5100, the port LED remains unlit. PR1317750
- On a QFX10016, permanent traffic loss is seen for some hosts after the initial ARP timer expiry caused by an ARP entry is not synchronized between the two PE devices. PR1322288
- On the QFX10016 EVPN-VXLAN scaled testbed, it takes up to 3 minutes for traffic to converge when a configuration related to a tenant (five IRB interfaces or VLANs) is added. PR1323042
- In a MH EVPN-VXLAN scenario, with IGMP snooping configured, in a scaled scenario: 1) For 10000 s,g scale: Trigger: disable DF link for convergence: Total convergence for 10000 s,g scale is 4.5 seconds with traffic rate of 60 kpps. Per flow convergence loss ranges from 3.16 seconds to 5.66 seconds 2) For 8000 s,g scale: Trigger :disable DF link for convergence: Total convergence for 8000 s,g scale is 2.86 seconds with traffic rate of 60 kpps. Per flow convergence loss ranges from 1.86 seconds to 3.73 seconds. PR1323155
- Traffic statistics for multicast stream on gr-interfaces does not work on a QFX5000 platform. PR1323622
- With 100-Gigabit DAC/copper cable connected between QFX5210-64C and QFX10000 devices, links might not come up reliably. The rest of the 100-Gigabit optics/AOC and 40-Gigabit optics/DAC/copper work well when connected between QFX5210-64C and QFX10000 devices. PR1324600
- Configuration of **mac-table-size** under VLAN switch options is not supported for QFX10002-60C. PR1325315
- In QFX5210-64C, irrespective of the physical interface speed, the speed displayed for gr-interface is always 800 mbps. PR1325695
- The **mac-learning-limit** option is not supported under VLAN switch options for the QFX10002-60C platform. PR1325752

- A few harmless error messages related to function rt_mesh_group_add_check() are seen during reboot. PR1335363
- Traffic statistics do not get updated on the gr-0/0/0 interface with ECMP. PR1335670
- On switching platforms, LACP aggregate Ethernet minimum-link with sync-reset enabled feature is not supported on an aggregated interface where micro-BFD is enabled. PR1342657
- Hardware watchdog does not work on QFX10008 and QFX10002-60C/PTX10002-60C platforms. PR1343131
- When the routes are changed from V4 to V6 or vice versa, routes are getting added from STC before all previous routes are deleted. Hence, error messages are seen. PR1350719
- The 100-Gigabit Ethernet interface goes down after Ethernet loopback is configured or deleted. PR1353734
- On the QFX5100, if a scaled configuration involving a LAG interface, more than 3000 VLANs, and corresponding next hops is removed and a new configuration involving a LAG interface is applied at the same time, the new configuration might not take effect until the previous configuration has been deleted. During this time, the FXPC process might utilize high CPU resources. PR1363896
- GRE tunnel next hop as ECMP is not supported. PR1368653
- On Junos OS Release 18.2R2, the intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface. PR1388082
- RE-ARP requests fail because they are sent without VLAN-ID. PR1390794

Interfaces and Chassis

- As the **link-speed** configuration statement cannot be hidden, unexpected behavior is observed with MC LAG peer status. PR1329030
- The supported ARP scale over MC-LAG interfaces is 48,000. PR1334321

Layer 2 Features

- On QFX5100 Virtual Chassis interfaces on which flexible VLAN tagging has been enabled, STP, RSTP, MSTP, and VSTP protocols are not supported. PR1075230
- In a QFX5210-64C platform, resilient hashing is not supported for LAG interfaces. PR1325499
- Packet statistics are not supported for logical child members of aggregated Ethernet interface. PR1335454
- **Targeted-broadcast forward-only** does not broadcast the traffic. PR1359031
- With **IGMP snooping** enabled on the leaf switches, multicast traffic is forwarded to VLAN/VNI, which does not have an active receiver. PR1388888

Routing Protocols

- The route unidimensional limit in Junos OS Release 18.1R1 is 1.6 million routes. PR1320865
- Removal and adding of em0 configuration cause physical interface to be reconfigured. This might cause BFD to flap if aggressive BFD timers are configured because of the hardware interrupt in the kernel. QFX5100 platform does not support BFD for minimum interval of less than 1 second. PR1332229

Virtual Chassis

- Virtual Chassis internal loop might happen at a node coming up from a reboot. During nonstop software upgrade (NSSU) on an QFX5100 Virtual Chassis, a minimal traffic disruption or traffic loop (greater than 2 seconds) might occur and it is considered to be known behavior. PR1347902

SEE ALSO

New and Changed Features | **283**

Changes in Behavior and Syntax | **291**

Known Issues | **300**

Resolved Issues | **308**

Documentation Updates | **326**

Migration, Upgrade, and Downgrade Instructions | **326**

Product Compatibility | **340**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 18.2R3.

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

EVPN

- In a scaled setup, if mac-move is triggered more than four times, the MAC move detection might not be reliable. PR1284315
- The chained-composite-next-hop (CNH) is a must for EVPN pure type 5 with VXLAN encapsulation. Without this Packet Forwarding Engine might not program the tunnel next hop. You have to explicit set it on QFX5110 using **set routing-options forwarding-table chained-composite-next-hop ingress evpn**, QFX10,000 it is applied as part of default configuration. **user@host> show configuration routing-options forwarding-table | display inheritance defaults**. PR1303246
- On QFX10000, in an EVPN collapsed L2 and L3 multihomed GWs topology, when traffic is sent from IP fabric toward EVPN, some traffic loss is seen. If the number of hosts behind EVPN gateways is increased, the traffic loss becomes higher. PR1311773
- In an EVPN-VXLAN scenario, ARP table information is not synchronized on two spines after reconfiguring an end host on a multihomed CE interface from IP1/MAC1 to IP1/MAC2. PR1330663
- When VTEP scale of more than 200 is used in Junos OS release 18.1R1, VTEPs might not come up for all the tunnels and might impact traffic. PR1342175
- On QFX5110 and QFX5200 switches that are configured to tunnel Q-in-Q traffic in an EVPN-VXLAN network, the pop operation does not work on ingress interfaces. PR1344102
- 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

- Certain QFX Series devices do not pad Ethernet packets with zeros. Therefore, some packets can contain fragments of system memory or data from previous packets. This issue is also known as Etherleak and often detected as CVE-2003-0001. For more information see https://kb.juniper.net/JSA10773. PR1063645
- 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 percent and continuous traffic drops are observed after bringing down/up the downstream interface or when an FPC comes online after FPC restart. This occur with multicast replication for 1000 VLAN/IRB's. PR1161485
- Because of the incorrect type defined when printing value drawn from Packet Forwarding Engine, the output for the show interfaces command might show an extremely large number instead of the actual value. For example, **user@host> show interfaces extensive et-0/1/0 Physical interface: et-0/1/0, Enabled, Physical link is Up . Input errors: Errors: 0, Drops: 0, Framing errors: 0, Runts: 0, Bucket drops: 0, Policed discards: 0, L3 incompletes: 0, L2 channel errors: 0, L2 mismatch timeouts: 0, FIFO errors: 0, Resource errors: 429496729600 <<<<<<<<<<<<<<< Output errors: Carrier transitions: 3, Errors: 0, Drops: 0, Collisions: 0, Aged packets: 0, FIFO errors: 0, HS link CRC errors: 0, MTU errors: 0, Resource errors: 0, Bucket drops: 0 Where 429496729600 is 6400000000 in hexadecimal**. However, the actual value is 0x64 or 100 in decimal. PR1245748
- Single-bit and multiple-bit ECC errors are not logged on QFX5110 switches. PR1251917
- On QFX10000 Series switches, at initialization, the port group module comes up after some time and negative ACKs are seen until the port group module is up. After the port group module is up, negative ACKs are no longer observed. This is an expected behavior because of an aggressive link scan feature introduced in Junos OS Release 17.2. PR1271579
- Traffic drop occurs because of CRC errors when traffic is sent over et interfaces. PR1313977
- Port LEDs on QFX5100 do not work. If a device connects to a port on QFX5100, the port LED stays unlit. PR1317750
- There might be traffic loss on the ingress PE device after a new EVPN neighbor is added or an existing EVPN neighbor is deleted. PR1319770
- Interface uptime has increased by 8 seconds from Junos OS Release 17.4R1 to Junos OS Release 18.1R1. Also, SDK upgrades across releases might impact the parameters such as login prompt appear time, FPC up time, and interface up time after switch reboot. PR1324374
- The management process (mgd) might panic after modifying aggregated Ethernet interface members under **ethernet-switching vlan** stanza. As a result, the remote session is terminated. PR1325736
- When user configures an L2 filter with mixed L2 and L3 match condition, error syslog is displayed to the user. As a workaround, commit is denied when mixed L2 and L3/L4 match conditions are configured on a L2 filter. PR1326715

- On QFX5200 standalone devices with VXLAN configured, the user-configured ingress ACL scale limit is 256 terms. PR1331730
- A BFD session over aggregated Ethernet flaps when a member link carrying the BFD Tx flaps. PR1333307
- Refrain from committing MTU changes for GRE and underlying interfaces in single commit. For any GRE interface MTU update follow the mentioned workaround. PR1335739
- On QFX10002, QFX10008, and QFX10016, ND is incorrectly working on IRB/Layer 3 interface with discard filter. PR1338067
- Changing MTU for GRE and underlying interfaces in single commit will be a caveat for the IPv4 GRE feature. Refrain from committing MTU changes for GRE and underlying interfaces in single commit. For any GRE interface MTU update follow the mentioned workaround. PR1339601
- The issue is specific to flexible VLAN-tagged interface and does not happen if the interface is in trunk mode with EVPN-VXLAN configuration. PR1345568
- Downgrade from TVP image to a non-TVP image is not supported. Upgrade from a non-TVP to a TVP image is supported. PR1345848
- QFX10000 platform drops the Aruba wireless access point (AP) heartbeat packets. As result, the Aruba wireless AP cannot work. PR1352805
- This issue observed only with 100 LR4 optics in the warm boot stage of VM's during unified ISSU process, flap is observed only on peer port. Recommend no to use 100G LR4 during unified ISSU. PR1353415
- The 100-Gigabit Ethernet interface goes down after you configure and delete the Ethernet loopback configuration. PR1353734
- While hot swapping 100G and 40G BiDi optics, it is recommended to give a gap of 4 to 5 seconds before you remove and re-insert. PR1356502
- On QFX5100 platforms with sFlow enabled, when deleting or deactivating the sFlow interface, all other interfaces might go down and fxpc core files are 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
- On QFX5210 switches, the filter with routing-instance applied to family inet logical interface causes traffic drop and gets discarded on unrelated interfaces. PR1364020
- From Junos OS Release 17.3R1, on the QFX10002 platform, in a rare condition, the IPFIX flow statistics (packet/byte counters) are incorrect in the exported record. Since the statistics are not collected properly, the flow might time out and get deleted because of the inactive timeout, causing the number of exported records to be sent out unexpected. Traffic spikes generated by IPFIX might be seen. PR1365864
- 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

- The statement **pm4x25_line_side_phymod_interfa** might throw the error **ERROR: u=0 p=81 interface type 16 not supported by internal SERDES for this speed 50000**. This error messages are seen when channelization is detected in the Junos OS Release 18.1R3. PR1366137
- On the QFX10000 line of switches, with EVPN-VXLAN, the following error message is seen: **expr_nh_fwd_get_egress_install_mask:nh type Indirect of nh_id: # is invalid**. PR1367121
- When any CLI command is executed immediately after the AIS script package is installed, then no output is generated. PR1368039
- The user might not be able to stop the ZTP bootstrap when a QFX10016 or QFX10008 router with more number of line cards is powered ON with the factory default configuration. PR1369959
- The L2 bridge domain might not be created on Packet Forwarding Engine after changing VLAN configuration. PR1371611
- Static speed of 100M setting remains after changing the speed 100M to auto-negotiation. PR1372647
- Beginning in Junos OS Release 17.1R1, the MAC address of the interfaces on the QFX10002-36Q and QFX10002-72Q will change. On the QFX10002-36Q, after the upgrading to Junos OS Release 17.x, the last octet of the interface MAC addresses increases by 3. On the QFX10002-72Q, after the upgrading to Junos OS Release 17.x, the last octet of the interface MAC addresses increases by 6. PR1375349
- In Junos OS Release 18.1R3, when one 50-Gigabit Ethernet port is taken down using the **ifconfig** command, the other port also goes down. PR1376389
- In certain scenario's where flows are sampled through aggregate bundles when jflow sampling is enabled, the following harmless error logs can be seen: **[Tue Oct 30 18:17:40.648 LOG: Info] expr_get_local_pfe_child_ifl: cannot find child ifl of agg ifl 74 for this fpc [Tue Oct 30 18:17:40.648 LOG: Info] flowtb_get_cpu_header_fields: Failed to find local child ifl for 74 [Tue Oct 30 18:17:40.648 LOG: Info] fpc0 cannot find stream on [hostname]**. PR1379227
- LOC and Diag system LEDs on the front panel are not defined yet. PR1380459
- In case multiple LLDP sensors are getting exported together and part of their keys are overlapped, data for these sensors can get skipped sometimes from being exported. PR1382691
- Last reboot reason is not correct if the device is rebooted because of power cycle. Last reboot reason will be displayed as vJunos OS reboot even if the device get rebooted because of the power cycling. PR1383693
- Port-mirroring-instance or analyzer-based mirroring does not work with input as VLAN ingress when VLAN is mapped to VXLAN. PR1384732
- On QFX10008 and QFX10016 platforms, 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
- On Junos OS Release 18.4R1, the intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface. PR1388082

- When the **show** command takes a long time to display results, the STP might change its status because the BPDUs are no longer processed and cause outages. PR1390330
- On QFX10000 switches, the major alarm **FPC Management Ethernet Link Down** might be displayed for the management Ethernet (em0 or em1) interface that is administratively down. The alarm message has no service impact and can be ignored. PR1391949
- If PTP transparent clock is configured on the QFX5200, and if **IGMP snooping** is configured for the same VLAN as PTP traffic, the PTP over Ethernet traffic might be dropped. PR1395186
- L2 multicast and broadcast convergence is high while deleting and adding back the scale configurations of VLANs and VXLAN. PR1399002
- On QFX5100, the traffic initiated from a server connected to an interface might 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
- On QFX10002, QFX10008, and QFX10016, a auto correctable non-fatal hardware error on PE chip (which is ASIC on QFX10002, the third-generation FPC on PTX3000/PTX5000, and the Line card on QFX10008/QFX10016) is reported as 'FATAL' error. Hence, the related Packet Forwarding Engine might 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
- When the storm control profile is applied on MC-AE interface, even if the traffic exceeds the bandwidth of the storm configuration does not shut down. Because traffic is not going through this policer-based rate limiting algorithm. PR1411338
- On QFX10,000 platforms with EVPN, if an EVPN instance is created through the statement **set protocols evpn encapsulation mpls**, then the MAC learning might not happen on the CE-facing interface if the interface is configured with trunk-mode, because EVPN/MPLS is not supported on QFX10000 Series devices. PR1416987
- On the QFX5110 platforms, uRPF check in strict mode might not work appropriately. PR1417546
- ERSPAN traffic is not tagged when the output interface is a trunk port. PR1418162
- On Junos OS routers and switches with Link aggregation control protocol (LACP) enabled, deactivating a remote aggregated Ethernet member link will make the local member link move to LACP detached state. The detached link will be invalidated from the Packet Forwarding Engine aggregated Ethernet-forwarding table as expected. However, if the device is rebooted with this state, all the member links will be enabled in Packet Forwarding Engine aggregated Ethernet-forwarding table irrespective of LACP states and result in traffic drop. PR1423707
- When channelization is configured on FPC QFX10000-30C (ULC-30Q28) while jFlow (jFlow v9 or v10) is configured on this board, the jFlow export might fail. As a result, loss of sample flow is seen. PR1423761
- CRC errors can be seen when other manufacturer device is connected to QFX10000 on a 100-gigabit link with QSFP-100GBASE-LR4-T2. Other manufacturer device report CRC errors and input errors on those 100-gigabit links. The QFX10000 interfaces do not show any errors causing packet loss. PR1427093

- On QFX10000 Series platforms, the range of Maximum Transmission Unit (MTU) allowed for Layer 2 interface is from 256 to 9216. However, when configuring MTU lower than 270 (256 to 269 inclusive), the Layer 2 traffic drop is seen because of the defective MTU check. PR1431902
- VRRP-V6 state is flapping with init and idle states after configuring **vlan-tagging**. PR1445370

Infrastructure

- FTP displays the following messages: **ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)**. PR1315605

Layer 2 Ethernet Services

- In an MC-LAG with force-up scenario, an LACP PDU loop might be seen when both MC-LAG nodes and the access device use the same admin key. PR1379022
- On QFX5100 and QFX5200 line of switches with spine-leaf scenario, when some (two or more than two) underlay interfaces with ECMP are brought down on a leaf device, the multihop BFD overlay sessions between spines and leafs might flap. And if BFD flaps, the protocols depending on BFD (typically, IBGP Protocol) might also flap, which leads to traffic impact. PR1416941

Layer 2 Features

- On QFX10016, after delete and re-adding of 1000 lag interfaces, traffic drops could be seen until ARP are refreshed even though all lag interfaces come up. PR1289546
- The **Targeted-broadcast forward-only** command does not broadcast the traffic. PR1359031
- On a QFX5100 Q-in-Q might stop working for certain **vlan-id-list** configured under a physical interface. As a result, a Packet Forwarding Engine binary issue is addressed through an upcoming image. PR1395312
- On Junos OS QFX5000, on the 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 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
- On QFX5110 devices, stale entries might fill up the L3 egress table, preventing new entries from being added. This might impact traffic. PR1423368
- On QFX5000 platforms, the fxpc might crash repeatedly when a firewall filter is applied on a logical unit of a DSC interface. This impacts traffic. PR1428350
- Firewall counters of VXLAN access ports might not show correct values after child members are deleted or added in aggregated Ethernet interfaces. PR1441424

MPLS

- There could be some lingering RSVP state that would keep some labeled routes programmed in the Packet Forwarding Engine longer than they should be. This RSVP state will eventually expire and then delete the RSVP MPLS routes from FIB. However, traffic loss is not anticipated because of this lingering state or the corresponding label 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 do not increment LSP statistics signaled by RSVP-TE. PR1362936

Platform and Infrastructure

- In configurations with IRB interfaces, when interfaces are deleted (for example, FPC reboot), the Packet Forwarding Engine might log errors stating **nh_ucast_change:291Referenced l2ifl not found**. This condition should be transient, with the system reconverging on the expected state. PR1054798

Routing Protocols

- For the QFX10002 and QFX10008 switches, you might observe an increase in the convergence time of OSPF routes when compared to Junos OS Release 17.3. An average increase of 1.5 seconds is seen for 100,000 OSPFv3 routes. PR1297541
- In a PVLAN configuration, the isolated VLAN and community VLAN should not use same VLAN ID. PR1323520
- 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 routeable lo0 address. Using loopback address reduces dependency with interfaces. Note: Multihop is enabled for iBGP peers by default. PR1323557
- The VLAN range shown in community VLAN is 1.4094. Hence, VLAN 0 should not be configured as community VLAN in PVLAN. PR1323719
- When MoFRR is enabled, the traffic statistics on the multicast route show double the outgoing traffic because accounting is done for both the primary and backup route. When one of the upstream interfaces goes down, this issue is not seen. There is no workaround for this issue. PR1326338
- Higher convergence time for LFA with BFD occur in Junos OS Release 18.1. PR1337412
- 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 the QFX Series switches (except for QFX10,000 line of switches, if host-destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log/syslog action (for example, 'filter <> term <> then log/syslog'), such packets might not be dropped and might reach the Routing Engine unexpectedly. PR1379718

- The B**RCM_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
- On QFX5110 and QFX5200 platforms, the dcpfe might crash if any interface flaps. PR1415297
- In BGP graceful restart scenario, including helper mode which is enabled by default, rpd might crash and generate a core file because of the improper handling of BGP graceful restart stale routes while deleting the BGP neighbor. The rpd might crash and service/traffic impact might occur. PR1427987

SEE ALSO

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.2R3

Authentication and Access Control

- Without configuring anything related to dot1x, the syslog **dot1xd[2192]: task_connect: task PNACAUTH./var/run/authd_control addr /var/run/authd_control: Connection refused** is generated repeatedly. PR1406965

Class of Service (CoS)

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

EVPN

- The rpd process might crash with EVPN type-3 route churn. PR1394803
- 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 the local configuration is changed. PR1396915
- In the non-collapsed (centralized) topology, when one of the 2 spines deactivates the underlay protocol (OSPF), the leaf still points the virtual-gw-mac's next hop to the down spine. PR1403524
- EVPN routes might show "Route Label: 0" in addition to the real label. PR1405695
- The rpd might crash after NSR switchover in an EVPN scenario. PR1408749
- ARP entry points to failed VTEP after PE-CE link fails for multihomed remote ESI. PR1420294
- Multicast MAC addresses are learned in the Ethernet switching table with VXLAN through an ARP packet in a pure L2 configuration. PR1420764
- The device might proxy the ARP probe packets in an EVPN environment. PR1427109

Forwarding and Sampling

- The kernel crash might be observed when there is a firewall filter modification. PR1365265
- On Junos OS, firewall filter terms named "internal-1" and "internal-2" are ignored. PR1394922

General Routing

- The 1-gigabit Ethernet copper module interface shows "Link-mode: Half-duplex" on QFX10000 line platforms. PR1286709
- SFP-T might not work on QFX5100 and QFX5110 devices. PR1366218
- Packet Forwarding Engine is in a bad state after performing optics insertion or removal on a port. PR1372041
- The backup member switch might fail to become the master switch after switchover on QFX5100 and QFX5200 Virtual Chassis platforms. PR1372521

- RIPv2 update packets might not be sent when **IGMP snooping** is enabled. PR1375332
- Packet Forwarding Engine might get wedged if there are interfaces going to the down state. PR1376366
- Debug log message, **expr_nh_flabel_check_overwrite: Caller nh_id params**, classified as an error log when it should be LOG_INFO. PR1377447
- The overlay ECMP might not work as expected on QFX5110 in an EVPN-VXLAN environment. PR1380084
- There is an inconsistency in applying scheduler map with excess-rate on the physical interface and aggregated Ethernet interface. PR1380294
- Traffic drops and get discarded by FPC offline in MC-LAG scenario. PR1381446
- The QFX-QSFP-40G-SR4 transceiver might not be recognized after upgrading Junos OS on QFX5100e. PR1381545
- Static default route with next-table inet.0 does not work. PR1383419
- The log of **RPD_KRT_Q_RETRIES: list nexthop ADD: No such file or directory** might be continuously shown after the rpd process restarts. PR1383426
- DMA failure errors might be seen when the cache is flushed or the cache is full. PR1383608
- The Virtual Chassis could not come up after upgrading to QFX5E platforms. PR1383876
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevents a major alarm. PR1384435
- The QFX10K-12C-DWDM line card might crash when booting up. PR1386400
- The rpd process might end up with KRT queue might get stuck in a VRF scenario. PR1386475
- QFX5100, QFX5110, QFX5200, and QFX5210 Virtual chassis could not be formed normally. PR1387730
- Certain log messages might be observed on QFX Series platforms. PR1388479
- On QFX5100 Virtual Chassis, ARP received on SP-Style interface are not sent to all RVTEPs. Normal BUM traffic works fine. PR1388811
- FPC might crash on QFX5100 platforms in a large scale scenario. PR1389872
- An incorrect error message might be seen when J-Flow sensors are configured with reporting rate less than 30 seconds. PR1390740
- 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 QFX10,000 devices. PR1393937
- Unable to install licenses automatically on QFX Series platforms. PR1395534
- The subscriber bindings might not be successful on QFX Series platforms. PR1396470
- On QFX5110 switches, the Fan LED turns amber randomly. PR1398349

- The interrupt process consumes high CPU because of the intr{swi4: clock (0)} on QFX5100-48t-6Q running a QFX5100 Series image and Junos OS Release 18.x code. PR1398632
- The DHCPv6 relay packets are dropped when both the UDP source and destination ports are 547. PR1399067
- CPU hog might be observed on QFX10,000 Series platform. PR1399369
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. PR1399683
- SFP-LX10 transceivers do not work on QFX5110. PR1399878
- PEM I2C failure alarm might be shown incorrectly as failed. PR1400380
- Only one Packet Forwarding Engine might be disabled on FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- The authd might crash when issuing **show network-access requests pending** command when restarting the authd process. PR1401249
- File permissions are changed for **/var/db/scripts** files after reboot. PR1402852
- The STP does not work when aggregated Ethernet interfaces number is "ae1000" or above in QFX5000 and "ae480" or above in other QFX Series switches. PR1403338
- The DHCP discover packets are forwarded out of an interface incorrectly if DHCP snooping is configured on that interface. PR1403528
- The VRRP VIP might not work when it is configured on the LAG interface. PR1404822
- ARP/ND is not resolved if native VLAN ID configured for LAG access interface. PR1404895
- Executing the **request system configuration rescue save** command might fail with error messages. PR1405189
- DHCP might not work for some clients in dual AD fusion setup on EP ports. PR1405495
- The DHCP discover packets might be dropped over VXLAN tunnel if DHCP relay is enabled for other VXLAN/VLANs. PR1408161
- MAC address movement might not happen in flexible Ethernet services mode when family inet/inet6 and vlan-bridge are configured on the same physical interface. PR1408230
- Fan failure alarms might be seen on QFX5100-96S after upgrading to Junos OS Release 17.3R1. PR1408380
- Restarting line card on QFX10008 and QFX10016 with MC-LAG enhanced-convergence might cause intra-VLAN traffic to get silently dropped and discarded. PR1409631
- The FPC might crash and might not come up if **interface-num** or next hop is set to maximum value under **vxlan-routing** on QFX Series platforms. PR1409949
- LLDP memory leak when ieee dcbx packet is received in autonegotiation mode followed by another dcbx packet with none of ieee_dcbx tlvs present. PR1410239

- On EX2300-24P, the following error message is observed. **dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find if family**. PR1410717
- Traffic loss might be observed after VXLAN configuration change. PR1411858
- The spfe on satellite device in a Junos Fusion setup might crash and it could cause the satellite device to go offline. PR1412279
- The PEM alarm for the backup FPC remains on the master FPC though backup FPC is detached from the Virtual Chassis. PR1412429
- Junos PCC might reject PCUpdate/PCCreate message if the metric type is not type 2. PR1412659
- On QFX5000, EVPN or VXLAN mutlicast next-hop limit is 4000. PR1414213
- Virtual Chassis ports using DAC might not establish links on the QFX5200. PR1414492
- VXLAN encapsulation nexthop (VENH) does not get installed during BGP flap or restart routing. PR1415450
- Traffic loss might be seen on the aggregated Ethernet interface on QFX10000 platforms. PR1418396
- Rebooting QFX5200-48Y using **request system reboot** does not take physical links offline immediately. PR1419465
- Ping fails over Type-5 tunnel on IRB interfaces under EVPN-VXLAN scenario. PR1420785
- Error messages might be seen on QFX10,000 platforms during DFE tuning. PR1421075
- BFD might get stuck in slow mode on QFX10002, QFX10008, and QFX100016 platform. PR1422789
- QFX5100-48T 10G interface might be autonegotiated at 1-Gbps speed instead of 10-Gbps. PR1422958
- The interface cannot start up when the remote-connected interface only supports only 100M in QFX5100 Virtual Chassis setup. PR1423171
- All interfaces might go down and the dcpfe might crash if SFP-T is inserted on QFX5210. PR1424090
- IPv6 neighbor solicitation packets for link-local address are dropped when passing through QFX10002-60C. PR1424244
- All interfaces creation failed after NSSU. PR1425716
- Heap memory leak might be seen on QFX10,000 platforms. PR1427090
- Licenses using the flag for OVSDB on **show system license** might not be flagged even though OVSDB is configured and working. PR1428207
- On EVPN-VXLAN, the L2ALD generates a core files are generated when number of VXLAN hardware IFBDS exceeds the maximum limit of 16,382. PR1428936
- DHCP-relay might not work in an EVPN-VXLAN scenario. PR1429536
- Interface on QFX Series devices does not come up after the transceiver is replaced with one having different speed. PR1430115

- On the QFX10000 line of devices, when incoming packets are processed by interfaces that have the hold-down timer configured, packets are forwarded through the ASIC. PR1430722
- On QFX switches **Validation of meta data files failed** on hypervisor. PR1431111
- On QFX5110 SFP-T, all ingress traffics are dropped on 100M fixed-speed port configured with no autonegotiation. PR1431885
- LASER TX remained enabled while interface is disabled using the Routing Engine CLI configuration. PR1436286
- Transit DHCPv6 packets might be dropped on QFX5100 and QFX5200 platforms. PR1436415
- On QFX5110, QFX5200, and QFX5210 switches, there is no jnxFruOK SNMP trap message when the power cable is disconnected and connected back. PR1437709

Interfaces and Chassis

- The dcpfe process might crash on using an unsupported GRE interface configuration. PR1369757
- Changing the value of **mac-table-size** to default might cause all FPCs to reboot. PR1386768
- The logical interfaces in EVPN routing instances might flap after committing configurations. PR1425339

Junos Fusion Satellite Software

- Extended port (EP) LAG might go down on the satellite devices (SDs) if the related cascade port (CP) links to an aggregation device (AD) go down. PR1397992

Layer 2 Ethernet Services

- The malfunction of the core isolation feature in EVPN-VXLAN scenarios causes traffic drop and gets discarded. PR1417729
- Continuous MAC change might cause CPU hogs and FPC reboots. PR1424653

Layer 2 Features

- VXLAN next-hop entry leak issue is seen on QFX5100 and QFX5200 line of switches. PR1387757
- With **IGMP snooping** enabled on the leaf switches, multicast traffic is forwarded to VLAN/VNI which does not have an active receiver. PR1388888
- On QFX Series switches, error message **Failed with error (-7) while deleting the trunk 1 on the device 0**. PR1393276
- On EVPN-VXLAN, the DCPFE restarts at the **_bcm_field_td_counter_last_hw_val_update** routine after upgrading spine with the latest image. PR1398251
- ARP response packets might include an incorrect VLAN ID and VNI. PR1400000
- The dcpfe process might crash when the Packet Forwarding System with scaled EVPN/VXLAN configuration restarts. PR1403305

- EVPN-VXLAN unicast IPv6 NS message gets flooded on Layer 3 gateway. Therefore, both IPv4 and IPv6 traffics get dropped on Layer 2 switch. PR1405814
- The IPv6 NS/NA packets received over VTEP from an ESI host are incorrectly flooded back to the host. PR1405820
- **IGMP-snooping** on EVPN-VXLAN might cause the flooding of OSPF hello packets after VTEP leaf reboot. PR1406502
- QFX5110 Virtual Chassis generates DDoS messages of different protocols on inserting a 1G/10G SFP or forming VCP connection. PR1410649
- With **arp-suppression** enabled, QFX5100 and QFX5200 might not forward IPv6 router solicitations or advertisements packets. PR1414496

Network Management and Monitoring

- The chassisd might crash and restart after the AGENTX session time out between master (snmpd) and sub-agent. PR1396967
- Log files might not get compressed during the upgrade. PR1414303

Platform and Infrastructure

- The **Platform failed to bind rewrite** message might be seen when chassis control is restarted with the CoS rewrite rule configured on aggregated Ethernet interface. PR1315437

Routing Protocols

- BUM packets might get looped if EVPN multihoming interface flaps. PR1387063
- Autonegotiation errors and flush operation failed error is seen after the power cycle of the device. PR1394866
- On QFX5110 and QFX5200 switches, EVPN-VXLAN non-collapsed dcfpe core file is seen at **brcm_pkt_tx_flush, l2alm_mac_ip_timer_handle_expiry_event_loc**, after the random event. PR1397205
- The rpd soft core file is seen and inappropriate route selection might be seen when Layer 2 VPN is used. PR1398685
- The FPC/dcpfe process might crash because of the interface flap. PR1408428
- Host-generated ICMPv6 RA packets might be dropped on the backup member of a Virtual Chassis if **igmp-snooping** is configured. PR1413543
- The QFX Series switches might not install all IRB MAC addresses in the initialization. PR1416025
- On QFX5200 switches, consistent traffic flow is seen. But the hash for ECMP next hop is not consistent. PR1422324
- After deleting IRB logical interface, MAC entry for the IRB is deleted for the IRB hardware address, packets destined to other IRB logical interfaces where MAC is not configured. PR1424284

Spanning Tree Protocols

- The l2cpd might crash if the **VSTP traceoptions** and **VSTP VLAN all** commands are configured. PR1407469

Resolved Issues: 18.2R2

Class of Service (CoS)

- The packets with destination-address 224.0.0.0/4 cannot be matched by loopback filter. PR1354377

EVPN

- The QFX10000 might drop if transited traffic comes from MPLS network to VXLAN/EVPN. PR1360159
- The l2ald core file is generated at l2ald_get_bd_client in qfx10k2: EVPN-VXLAN. PR1365254
- Increased risk of routing crash with temporary impact on traffic on QFX10000 or QFX5100 nodes with certain configuration changes or clearing L2 or L3 learning information a high-scale EVPN-VXLAN configuration environment. PR1365257
- Proxy ARP might not work as expected in an EVPN environment. PR1368911
- QFX10000 or import default IPv6 route to VRF causes infinite entries to get created in 'evpn ip-prefix-database' and become unstable. PR1369166
- VTEP's MAC address might not be learned in the Ethernet switching table. PR1371995
- The statement **[evpn_vxlan] [virtual_switch] show ethernet-switching vxlan-tunnel-end-point esi** shows large number of MAC count on QFX10000. PR1394982

General Routing

- After zeroizing, QFX5100 is treating 40G AOC uplink as 4x10g breakout on enabling auto-channelization. PR1317872
- On the QFX10016 EVPN-VXLAN, it takes upto 3 minutes for traffic to converge when configured. PR1323042
- Port 0 does not come up in QFX5100-48t member in mixed VCF. PR1323323
- CoS is incorrectly applied on Packet Forwarding Engine leading to egress traffic drop. PR1329141
- Status LED on the chassis does not show up on QFX10002-60c. PR1332991
- AI-script does not get auto upgrade unless it is manually done after a Junos OS upgrade. PR1337028
- On QFX5100 platforms, LR4 QSFP can take up to 15 minutes to come up after Virtual Chassis reboot. PR1337340
- On the QFX10000 platforms, VRRP function does not work well when it is configured on sub-interfaces. PR1338256
- On QFX5200, unified ISSU from Junos OS Release 17.2X75-D41 to Junos OS Release 18.2 will be aborted when dcpfe crashes. PR1338300

- PAFXPC core file is generated when a remote member physical interface is referenced in **show dcbcm ifd <ifd-name>** on QFX5100 platform configured in a VC. PR1343701
- On QFX5100, FAN RPM fluctuates when temperature sensor reaches its threshold. PR1345181
- On QFX10000 platforms, NETCONF SSH TCP port 830 traffic hits host path or unclassified queue. PR1345744
- Backup Routing Engine might experience a crash, causing vmcore to be generated on master Routing Engine, master Routing Engine performance will not be affected. PR1346218
- Blackholing traffic with destination MAC matching the virtual gateway MAC might be seen. PR1348659
- 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
- Bogus DDoS counter values and syslog messages could be seen after clearing DDoS statistics for a specific protocol on QFX10000 Series switches. PR1351212
- ARP learning might fail after changing the interface MAC address. PR1353241
- On 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 IGMP membership report packets might not be forwarded over an interface on QFX10000. PR1360137
- On QFX10000 platform, packets will be dropped when virtual-gateway-address is configured on an IRB interface associated with a non-vxlan VLAN. PR1360646
- FEC is incorrectly displayed on QFX10002-36Q and QFX5110. PR1360948
- The GTP traffic might not be hashed correctly on aggregated Ethernet interface. PR1361379
- On QFX10000 platforms, the **clear services accounting statistics inline-jflow fpc-slot** command does not work. PR1362396
- VME interface might be unreachable after the link flap of em0 on master FPC. PR1362437
- Traffic might not be forwarded when the member link of the aggregated Ethernet is added or deleted. PR1362653
- 1G interface might stop working when "auto-negotiation" is off by default. PR1362977
- The kernel displays syslog message for the configuration **tcp_timer_keep Dropping socket connection**. PR1363186
- On QFX10008 and QFX10016 platforms, MPLS exp rewrite might not work for IPv6 and IPv4 traffic. PR1364391
- Traffic loss is observed when unified ISSU is performed when aggregated Ethernet interfaces is configured with LACP protocol. PR1365316
- Root password recovery process does not work. PR1365740

- The l2cpd process might crash when configuring MVRP with private VLAN and RSTP interface all. PR1365937
- SFP-T might not work on QFX5100 and QFX5110 devices. PR1366218
- The tagged traffic is dropped in the untagged EVPN/VXLAN scenario. PR1366336
- The chassisd might crash after issuing the CLI **show chassis hardware**. PR1366746
- On QFX10002-60C and QFX10000-30C platforms, some interfaces do not come up during initialization after a reboot. PR1368203
- On QFX Series switches IS-IS adjacency with another vendor's switches might go down. PR1368913
- The **commit** or **commit check** might fail because of the error **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- In certain routing topologies with sFlow configured, sampled packets might be duplicated and sFlow records are not sent to the collector. PR1370464
- On QFX10000 platforms, before the Junos OS Release 17.3R3 code, the maximum number of ESI logical interfaces is 4000 in the Packet Forwarding Engine. PR1371414
- On QFX5100, the IPv6 routed packet will be transmitted though VRRP state is in transition to master. PR1372163
- Packets might be dropped after deleting a filter from an interface. PR1372957
- MAC refresh packet might not be sent out from the new primary link after RTG failover. PR1372999
- TPI-50840 BUM traffic received on 5110 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 QFX10000 and non-Juniper device. PR1374321
- On QFX5110 Ethernet-switching flood group shows incorrect information. PR1374436
- Only loopback interface is supported under VRF routing instances. PR1375130
- Same address family [Subnet IFL or IRB IFL but not both] needs to be configured for establishing VTEPs. PR1376996
- The auto-negotiation interface might go down if the opposite device supports only 10/100M auto-negotiation. PR1377298
- Deleting an IRB interface might affect other IRB interfaces if the same custom MAC address is configured. PR1379002
- L3VPN traffic might be dropped because of one core-facing interface down. PR1380783
- A QFX5xxx packet forwarding engine (PFE) might show DISCARD next hop for **overlay-bgp-lo0-ip** in a leave-spine topology. PR1380795
- Virtual Chassis master is copying **/var/db/ovs** database to backup every 10 seconds which causes a high write IO and shorten the SSD lifetime in Open vSwitch Database (OVSDB) environment. PR1381888

- EVPN-VXLAN ARP and 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 Layer3 interface might stop pinging directly connected link address after deleting Layer2 on physical interface. PR1384144
- BFD sessions might flap consistently. PR1384601
- All 1G SFP copper and 1G fiber optic links remain up on QFX10008 after all SIBs/FPCs are offline. PR1385062
- The IPv6 packet might not be routed when IPv6 packet is encapsulated over IPv4 GRE tunnel on QFX10000. PR1385723
- The spine EVPN routes might get stuck in a hidden state with next hop as unusable after FPC is offline in the spine. PR1386147
- Intra PoD traffic drop observed with trap code sw.egnh.cfg_discard and VXLAN/VTEP programming missing. PR1387593
- CPSM daemon memory leak in VMHOST. PR1387903
- On QFX10000 platforms, MAC learning might stop working on some LAG interfaces after frequent MAC moves. PR1389411
- BFD flaps are seen on QFX10000 platforms with inline BFD. PR1389569
- IPv6 next hop programming issue is observed on QFX10016 device running on Junos OS Release 15.1X53-D67. PR1393937
- The l2ald core file seen when Layer 2 learning traceoptions are enabled. PR1394380
- On QFX5110 Virtual Chassis, after Routing Engine switchover, LACP will be down on a peer device and will never been recovered automatically. PR1395943
- If GRES and NSR is enabled on a QFX5100 (single Routing Engine), DHCP subscribers fails to bind. PR1396470

Infrastructure

- On QFX5100, enabling **mac-move-limit** stops ping on flexible-vlan-tagging enabled interface. PR1357742

Interfaces and Chassis

- MC-LAG peer does not send ARP request to the host. PR1360216
- On QFX products, the CLI allows to configure more sub-interface than the limit of 2048 sub-interfaces on lag interface from Junos OS Release 17.2R1. PR1361689
- On QFX5200 MCLAG, **parse_remove_ifl_from_routing_inst() ERROR: No route inst on et-0/0/16.16386**, errors are seen after restarting l2cpd daemon. PR1373927

Layer 2 Features

- On QFX5100, the Junos OS Release 14.1X53-D46.7 the storm control profile is missing for interfaces in hardware. 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 a huge memory is allocated. 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 VXLAN scenario. PR1368258
- When **native-vlan-id** is configured for aggregated Ethernet LACP session to multihomed server goes down. PR1369424
- A port might still work if it is deleted from an aggregated Ethernet interface. PR1372577
- DHCP discover packets might be dropped if there is VXLAN configured. PR1377521
- Packets might be dropped on AD in Junos Fusion Data Center environment. PR1377841
- The dcpfe process might crash while changing MTU of physical ports for GRE. PR1384517
- The LACP might be detached state when deleting **native-vlan-id** on aggregated Ethernet interface with **flexible-vlan-tagging** configured. PR1385409
- The dcpfe core might be observed when doing “restart routing" or BGP neighbors flaps when EVPN-TYPE 5 routes are present. PR1387360
- On QFX5000 switches, EVPN-VXLAN fails to forward the IPv6 NS packet from remote VTEP to local host. PR1387519
- The dcpfe process might crash after VXLAN overlay ping. PR1388103
- RTG MAC refresh packets will be sent out from non-RTG ports if the RTG interface belonging to Virtual Chassis master flaps. PR1389695
- Cisco Discovery Protocol (CDP) packets are not forwarded by QFX10000. PR1389829

MPLS

- LSP might not be established properly between QFX5000 and other devices. PR1351055
- The LSP might remain UP even if no path is acceptable because of CSPF failure. PR1365653
- NO-propogate-TTL acts on MPLS swap operation. PR1366804
- LSP with auto-bandwidth enabled goes down during HMC error condition. PR1374102
- On QFX10000, the LSP statistics and autobandwidth functionality do not work on QFX10002 with single hop LSP. It works with multi-hop LSP. PR1390445

Platform and Infrastructure

- On Junos OS, the next hop index allocation fails. The private index space exhausts through incoming ARP requests to management interface (CVE-2018-0063). PR1360039
- When migrating from VPLS to EVPN vlan-aware, after adding **routing-instance** configuration with **protocols evpn extended-vlan-list**, the traffic is dropped on Packet Forwarding Engine as "invalid L2 token". PR1368802
- Traffic is silently dropped and discarded with indirect next hop and load balancing. PR1376057
- LSI binding is missing upon nd6 entry refresh after l2ifl flap. PR1380590
- IRB interface does not turn down when master of Virtual Chassis is rebooted or stopped. PR1381272

Routing Protocols

- On QFX5100 platforms, parity errors in L3 IPv4 table in the Packet Forwarding Engine memory might result in traffic getting silently dropped and discarded. PR1364657
- The dcpfe might crash and all interfaces flap. PR1369011
- If a QFX5100 device has a host route with ECMP next hops and receives a better path with single next hop then the next hop in hardware will not be changed. PR1387713

Software Installation and Upgrade

- Commit might fail in single-user mode. PR1368986

Resolved Issues: 18.2R1

EVPN

- Error message **JPRDS_DLT_ALPHA KHT** shows as failed, but the entries in hardware are programmed correctly. PR1258933
- In an EVPN-VXLAN setup, IPv6 packet loss is observed after normal traffic run rate. PR1267830
- The sub interface from same physical port do not work if configured under 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. PR1287557
- VXLAN traffic loss is observed after deleting and adding VLANs. PR1318045
- Core link flap might result in an inconsistent global MAC count. PR1328956
- On QFX5100, with EVPN-VXLAN, the leaf device is forwarding traffic to the incorrect VTEP after MAC move/vmotion. PR1335431
- Traffic might be lost on Layer2 and Layer3 spine node in multihome EVPN scenario. PR1355165

- In an EVPN-VXLAN environment, BFD flap causes VTEP to flap and the Packet Forwarding Engine crashes. PR1339084
- The routing protocol process (rpd) crashes and generates a core file on QFX Series switches with multiple VLANs with vlan-id zero, unique VNID. PR1342351
- The traffic might get dropped because the core is down. PR1343515

General Routing

- C0 fiber link does not come up. PR1298876
- Traffic loss might be seen while sending traffic through the 40G interface. PR1309613
- Traffic loss is observed while performing NSSU. PR1311977
- Certain IGMP join packets cannot be processed correctly at a high rate. PR1314382
- Transit traffic over GRE tunnel might hit the CPU and trigger a DDoS violation on the L3 next hop. PR1315773
- Packets such as TDLS without IP header are looped between virtual gateways. PR1318382
- Chassis MIB SNMP OIDs for VC-B member chassis are not available after MX Series Virtual Chassis unified ISSU. PR1320370
- The MAC address get stuck with "DR" flag on the spine node even though packets are received on the interface from source MAC. PR1320724
- The OpenFlow session cannot be established correctly with controller and interfaces options configured on QFX5100 switches. PR1323273
- On a QFX10000 platform deployed in a spine layer without any CE interfaces attached, the ARPs will not get resolved on the spine, and traffic drop might be observed. PR1324739
- The GRE traffic is not decapsulated by the firewall filter. PR1325104
- VLAN or VLAN bridge might not be added or deleted if there is an IFBD HW token limit exhaustion. PR1325217
- Unable to configure persistent learning using CLI **set switch-options interface <interface-name>** because no option is found PR1325313
- MAC move is not expected when disabled globally with CLI **set protocols l2-learning global-mac-move disable-action**. PR1325524
- MAC aging is not happening on lag interface. PR1325555
- ARP request packets might not be flooded on QFX5110. PR1326022
- On QFX5210, when the physical interface is down, the CLI **show chassis LED** still shows "Green". PR1326078
- The major alarm about **Fan and PSU Airflow direction mismatch** might be seen by removing the management cable. PR1327561

- Deleting one VXLAN might cause traffic loop on another VXLAN in a multihoming EVPN and VXLAN scenario with the service provider style interface. PR1327978
- On QFX10002, a major alarm should be cleared once the chassis has more PEM units installed than the "minimum PEM" configuration. PR1327999
- A FAN tray removal or insertion trap is not generated for the backup FPC. PR1329031
- IRB physical interface static MAC address is not taking effect. PR1329032
- The CLI command **set chassis fpc 0 pic** has an option of PIC numbers 0 to 2, but the hardware only has one PIC. PR1329105
- The etherStatsCRCAlignErrors port counters might disappear in the SNMP tree. PR1329713
- After commit, members of Virtual Chassis or VCF are split and some members might get disconnected. PR1330132
- The rpd generates a core file on new backup Routing-Engine at **task_quit,task_terminate_timer_callback,task_timer_dispatch,task_scheduler after disabling NSR+GRES**. PR1330750
- On QFX10002-36Q, DHCP relay or server not working on GRE interface. PR1331158
- PTP BC with its PTP slave interface configured on a 100-Gigabit Ethernet interface might get stuck in FREERUN state. PR1331752
- Adding or deleting a tunnel configuration might result in FPC crash in a scaled GRE tunnels scenario. PR1331983
- On QFX5210, for some of the UFT profiles, the UFT is not able to scale the s,g entries to around 95 percent of the supported scale. PR1332170
- The error messages **out of HMC range** and **HMC READ faild** are seen. PR1332251
- Traffic does not flow through VCP ports after rebooting the Virtual Chassis members. PR1332515
- In an EVPN-VXLAN environment, DF drops multicast traffic. PR1333069
- The SDHCPv6 SOLICIT message is dropped. PR1334680
- Ethernet frame with Ethernet type of 0x8922 might be modified at egress by QFX10000. PR1334711
- The chassis reboots continuously when USB drive is connected after image recovery through USB and after CLI image install. PR1335269
- The supported scale for logical interface-based GRE tunnel on QFX10002-60C is 512. PR1335681
- The CLI command for beacon port state is not supported on QFX10002-60C. PR1337125
- SNMP jnxBoxDescr oid returns different value when upgrading to Junos OS Release 17.2. PR1337798
- The traffic coming from the remote VTEP PE device might get dropped. PR1338532
- The analyzer status might show as down when port mirroring is configured to mirror packets from an aggregated Ethernet member. PR1338564

- The VXLAN traffic might not be transmitted correctly with IRB interface as underlay interface of VTEP tunnel. PR1338586
- Reduced multicast scale with downstream IRB interfaces with snooping enabled. PR1340003
- Inconsistent result is seen in QFX5200 after using **deactivate xxx** command in pfc-priority and no-loss context. PR1340012
- IPv4 traffic routed out through incorrect interface after rpd restarts in leaf of IPCLOS profile. PR1341381
- In an EVPN-VXLAN, L3 traffic is not getting converged properly upon disabling the ECMP link between the spine and the leaf with EVPN-VXLAN configurations. PR1343172
- BPDU packets might get dropped and **bpdu-block-on-edge** might not work. PR1343330
- Broadcast frames might be modified with the ethertype 0x8850. PR1343575
- In an EVPN/VXLAN, VLAN with flexible-tag mode, the xe statistics is not updated for ingress. PR1343746
- Implement **[edit interfaces interface-name ether-options] configured-flow-control** option for QFX Series switches. PR1343917
- EVPN-VXLAN: ARP packet uses VRRP/virtual-gateway MAC in Ethernet header instead of IRB MAC address. PR1344990
- QFX5100 - Fan RPM fluctuates when temperature sensor reaches its threshold. PR1345181
- FXPC process might generate a core file while removing a VXLAN configuration. PR1345231
- Incorrect inner VLAN tag is sent from QFX10000 platform with Q-in-Q configured on the Layer3 sub-interface. PR1346371
- In QFX10000 SFlow scaling scenario, error messages are seen in syslog messages with respect to SFlow after configuring multiple LAG interfaces under SFlow protocol. PR1346493
- On QFX5100, in an EVPN a DCPFE core file is generated at **src/pfe/common/pfe-arch/brcm/applications/virtual/brcm_vxlan.c:2185**. PR1346980
- QFX5100-48T 10G interface might be auto-negotiated at 100M speed instead of 10G. PR1347144
- The IPFIX flow statistics are incorrect in the exported record. PR1347229
- Part numbers and serial numbers are not displayed for any of the 10G optics or DAC connected. PR1347634
- QFX10000 systems might encounter a chassis alarm indicating **FPC 0 Major Errors - PE Error code: 0x2100ba**. PR1347805
- Once in QFX10002-60C VMHOST crash is observed at prds_if_ifl_get_gre_stats (ifl=0x9288a608, expr_ifl_l2d_stats=0x2cd3790c), just after configuring GR interface on it. PR1348932
- The pfed process consumes 80 to 90 percent CPU running subscriber management on PPC-based routers. PR1351203
- DCPFE process might crash on QFX10000 switches. PR1351503

- The GTP traffic might not be hashed correctly for AE interface. PR1351518
- RPC output not showing failure when running **request system software add** with software already staged. PR1353466
- SFP-LX10 stay in up or down when connected. PR1353677
- The alarm errors might be seen during the bootup on QFX10000. PR1354582
- Untagged packets might not be forwarded through the trunk port. PR1355338
- On QFX5110 platforms, LX10 SFP needs to be reinserted after autonegotiation is enabled or disabled. PR1355746
- TPI-50840: qfx5110 ethernet-switching flood group shows incorrect information PR1374436
- Only loopback interface is supported under vrf routing instances PR1375130

Interfaces and Chassis

- If customer virtual local area network (CVLAN) range-16 (for example, vlan-id-list 30-45) is configured in a Q-in-Q (802.1ad) scenario, all the 16 VLANs might not pass traffic. PR1345994

Junos Fusion Satellite Software

- AD failure (power off) in a DC fusion is causing complete or partial traffic loss for extended period. PR1352167

Layer 2 Features

- MAC learning might fail for device on extended port of satellite device after MAC moving in a Junos Fusion scenario. PR1324579
- The DHCP discover packets might be looped in an MC-LAG and DHCP-Relay scenario. PR1325425
- In QFX5100, with multiple logical units configured on an interface, **input-vlan-map POP** does not remove outer vlan-tag when Q-in-Q and VXLAN are involved. PR1331722
- Push is not working for VXLAN local switching with the Q-in-Q. PR1332346
- Interface with **flexible-vlan-tagging** and **family ethernet-switching** does not work on QFX10000. PR1337311
- The DCPFE/FXPC process might crash and generate a core file. PR1362332

MPLS

- In a QFX5100, a unified ISSU is not supported with MPLS configuration. PR1264786
- A traffic drop is seen during NSR switchover for RSVP P2MP provider tunnels used by MVPN. PR1293014
- MPLS forwarding might not happen properly for some LSPs. PR1319379
- The rpd might crash on the backup Routing Engine because of memory exhaustion. PR1328974
- The hot standby for l2 circuit does not work on QFX5000. PR1329720

Multicast

- An aggregated Ethernet or IRB configuration causes kernel crash vmcore , and causes chassis or FPC reboot. PR1335904

Platform and Infrastructure

- The ARP might not update, and packets might get dropped at the Routing Engine. PR1348029
- When a Junos OS image is shipped with translation scripts downgrading to another image, stale symlinks of translation scripts at the time of mgd initialization leads box going into amnesiac state. PR1341650

Routing Protocols

- The **copy-tos-to-outer-ip-header** command is not supported, because of the hardware limitation. PR1313311
- Some of the IPv4 multicast routes in the Packet Forwarding Engine might fail to install and update. PR1320723
- In QFX5100, consistent hashing is not getting programmed. PR1322299
- QFX10002-60C is not supported as FHR in multicast PIM SM based network. PR1324116
- IS-IS L2 Hello packets are dropped when they come from a Brocade device. PR1325436
- Degradation is seen in some OSPF parameters and some of the RIB parameters are improved. PR1329921
- The loopbacked IRB interface is not accessible to the remote network. PR1333019
- The dcpfe crashes in a route leak scenario on QFX10000. PR1334714
- The rpf-check-policy does not work as expected. PR1336909
- On QFX5000 Series switches, BGP might be down due to the congestion state of CPU on receiving Ethernet pause frames. PR1343597
- DF is not working; ping fails if MTU is different on the interfaces. PR1345495
- The vrf-fallback on QFX5000 is not supported in ALPM mode. PR1345501
- IPv6 packets with hop-by-hop header cannot be matched using filters. PR1346052

SEE ALSO

New and Changed Features | **283**

Changes in Behavior and Syntax | **291**

Known Behavior | **297**

Known Issues | **300**

Documentation Updates | **326**

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 documentation for the QFX Series.

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. 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 **18.2** 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 18.2 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-x86-64-18.2-R2.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 18.2 **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-x86-64-18.2R3.9.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-x86-64-18.2R3.9.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 18.2R3.

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-x86-64-18.2R3.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-x86-64-18.2R3.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-x86-64-18.2R3.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-x86-64-18.2R3.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 337
- Upgrading the Software Using Unified ISSU on page 337

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-x86-64-18.2R3.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-x86-64-18.2R3.n-secure-signed.tgz ...
     Install jinstall-host-qfx-5-f-x86-64-18.2R3.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 | 283
- Changes in Behavior and Syntax | 291
- Known Behavior | 297
- Known Issues | 300
- Resolved Issues | 308
- Documentation Updates | 326
- Product Compatibility | 340

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 340

Hardware Compatibility

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

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

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 18.2R3 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.

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 18.2R3 for the SRX Series devices.

Release 18.2R3-New and Changed Features

There are no new features in Junos OS Release 18.2R3 for the SRX Series devices.

Release 18.2R2-New and Changed Features

Flow-Based and Packet-Based Processing

- **SRX5K-SPC3 card with flow support in chassis cluster mode (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.2R2, the SRX5K-SPC3 and SRX5K-SPC-4-15-320 (SPC2) cards can operate together in a mixed-mode configuration on the SRX5000 line of devices using the same slot number in both nodes. If you are adding the SPC3 SPCs to the SRX5000 devices, you must install the new SPCs in the lowest-numbered slot of any SPC that provides central point functionality. SPC3 interoperates with the SRX5000 I/O cards (IOC2, IOC3), Switch Control Boards (SCB2, SCB3), Routing Engines, and SPC2 cards.

  [See Understanding Flow support on SRX5K-SPC3 Platforms.]

VPN

- **PowerMode IPsec (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.2R2, PowerMode IPsec (PMI) is a new mode of operation that provides IPsec performance improvements using Vector Packet Processing (VPP) and Intel AES-NI instructions. PMI utilizes a small software block inside the Packet Forwarding Engine that bypasses the regular flow processing and take an express flow data path, and utilizes the AES-NI instruction set for optimized performance of IPsec processing.

  You can enable PMI processing by using the **set security flow power-mode-ipsec** command.

  The following features are supported with PMI:

  - Auto Discovery VPN (ADVPN)
  - Internet Key Exchange (IKE) functionality
  - AutoVPN
  - High-availability
  - IPv6
  - Stateful firewall
  - st0 interface
  - Traffic selectors

  [See Understanding PowerMode IPsec.]

- **SRX5K-SPC-4-15-320 (SPC2) and SRX5K-SPC3 (SPC3) support for IPsec VPN (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.2R2, all IPsec VPN features that were previously supported only on SPC3 (model number: SRX5K-SPC3) are now supported on both SPC2 (model number: SRX5K-SPC-4-15-320) and SPC3 installed in the SRX5000 line of devices operating in chassis cluster mode or in standalone mode.

  [See Understanding VPN Support for Inserting Services Processing Cards.]

Release 18.2R1-S3 New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 18.2R1-S3 for the SRX Series devices.

Junos OS Release 18.2R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/4200, 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-D130. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D130 are not available in 18.2 releases.

Junos OS Release 18.2R1-S1 and Junos OS Release 18.2R1-S3 supports SRX5K-SPC3. Junos OS for SRX Series documentation includes information about SRX5K-SPC3.

New features for security platforms in Junos OS Release 18.2R1-S3 include:

VPN

- **IPsec VPN support on SRX5K-SPC3 card (SRX5400, SRX5600, SRX5800)**—Starting in Junos OS Release 18.2R1-S3, SPC3 card supports IPsec VPN with AutoVPN networks in point-to-point secure tunnel mode with multiple traffic selectors, Dead peer detection (DPD), IKE fragmentation, and Site-to-site VPN (responder only).

  For SPC3 cards, you can only verify the tunnel mapping on different SPUs using the **show security ipsec tunnel-distribution** command. You can continue to use **show security ike tunnel-map** command to view the tunnel mapping on different SPUs with SPC2.

  The **show security ipsec tunnel-events-statistics** command is not supported on SPC3 card.

  [See show security ipsec security-associations.]

Release 18.2R1-S1 New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 18.2R1-S1 for the SRX Series devices.

Junos OS Release 18.2R1 supports the following Juniper Networks security platforms: vSRX, SRX300/320, SRX340/345, SRX550HM, SRX1500, SRX4100/4200, 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-D130. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D130 are not available in 18.2 releases.

Junos OS Release 18.2R1-S1 supports SRX5K-SPC3. Junos OS for SRX Series documentation includes information about SRX5K-SPC3.

New features for security platforms in Junos OS Release 18.2R1-S1 include:

Hardware

- **SRX5K-SPC3 Card support (SRX5400, SRX5600, SRX5800)**—Starting with Junos OS Release 18.2R1-S1, SRX5K-SPC3 Services Processing Cards (SPCs) are available on SRX5400, SRX5600, and SRX5800 Services Gateways. SRX5K-SPC3 card provides additional processing power to run integrated services such as firewall, IPSec, and IDP. The SRX5K-SPC3 contains two Services Processing Units (SPUs) with 128GB of memory per SPU. All traffic traversing the services gateway is intelligently distributed by I/O cards (IOCs) to the SPUs to have services processing applied to it.

  [See SRX5400 Services Gateway Hardware Guide, SRX5600 Services Gateway Hardware Guide, SRX5800 Services Gateway Hardware Guide, and SRX5400, SRX5600, and SRX5800 Services Gateway Card Reference.]

Interfaces and Chassis

- **User visibility improvements for chassis environment CLI (SRX5400, SRX5600, SRX5800)**—Starting in Junos OS Release 18.2R1-S1, the **show chassis environment fpc** CLI command displays current and power for SPC3 board along with the FPC voltage. In the earlier releases, only FPC voltage was displayed.

  [See show chassis environment.]

J-Web

- **J-Web supports SRX5K-SPC3 Card**—Starting Junos OS Release 18.2R1-S1, J-Web is enhanced to show SRX5K-SPC3 card support for SRX5400, SRX5600, and SRX5800 devices.

Platform and Infrastructure

- **SRX5K-SPC3 card (SRX5400, SRX5600, SRX5800)**—Starting in Junos OS Release 18.2R1-S1, a new service processing card (SRX5K-SPC3) is introduced for the SRX5000 line of devices. The introduction of the new card improves the scalability and performance of the device and maintains its reliability as it preserves the chassis cluster functionality. The SRX5K-SPC3 card supports higher bandwidth for service processing. It provides support for the following software features:
  - Application layer gateway (ALG)
  - Advanced anti-malware (Juniper Sky ATP)
  - Application security suite
  - Flow-based packet processing implementation
  - GPRS tunneling protocol (GTP) and stream control transmission protocol (SCTP)
  - High availability (chassis cluster)
  - Intrusion detection and prevention (IDP)
  - J-Web
  - Network address translation (NAT)
  - Stateful firewall

- SSL proxy
- Firewall user authentication
- UTM (antivirus, web filtering, content filtering, and antispam)

> **NOTE:**
> The following limitations apply for the SPC3 card in Junos OS Release 18.2R1-S1:
>
> - Interoperability of SPC2 card and SPC3 card is not supported.
> - IPsec VPN functionality is not supported with SPC3 card.

[See Understanding Flow support on SRX5K-SPC3 Platforms,Monitoring of Global-Level Objects in a Chassis Cluster, and Persistent NAT and NAT64.]

Release 18.2R1 New and Changed Features

This section describes the new features and enhancements to existing features in Junos OS Release 18.2R1 for the SRX Series devices.

ALG

- **TWAMP supports for ALG traffic (SRX Series)**—Starting in Junos OS Release 18.2R1, the Two-Way Active Measurement Protocol (TWAMP) Application Layer Gateway (ALG) is supported to enable the TWAMP data traffic to pass through the SRX Series device without needing a predefined policy permission.

  [See Understanding the Two-Way Active Measurement Protocol (TWAMP) Application Layer Gateway (ALG).]

Application Security

- **Application Quality of Experience (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100 and SRX4200, vSRX)**—Starting in Junos OS Release 18.2R1, AppQoE enables you to effectively prioritize, segregate, and route business-critical applications traffic without compromising performance or availability.

  AppQoE utilizes the capability of application identification and advanced policy-based routing to identify specific applications in the network and to specify a path for the application traffic according (service-level agreement) SLA rules.

  AppQoE monitors RTT, jitter, and packet loss on each link, and based on the score, seamlessly diverts applications to an alternate path if the performance of the primary link is below acceptable levels as specified by the SLA. Measurement and monitoring of application performance is done using active and passive probes, which detect SLA violations and help select an alternate path for that particular application.

  [SeeApplication Quality of Experience.]

- **Support for advanced policy-based routing (APBR) policy (SRX Series, vSRX)**—Starting in Junos OS Release 18.2R1, you can configure advanced policy-based routing (APBR) policies by defining source addresses, destination addresses, and applications as match conditions; and after a successful match, the configured APBR profile is applied as an application service for the session.

  In previous releases of Junos OS, an APBR profile could be attached to an incoming security zone of the ingress traffic, and the APBR was applied only on the basis of the security zone.

  This enhancement provides more flexible traffic-handling capabilities that offer granular control for forwarding packets.

  [See Advanced Policy-Based Routing.]

Authentication and Access Control

- **Support for user firewall to configure ClearPass and JIMS at the same time (SRX Series, vSRX)**—Starting in Junos OS Release 18.2R1, you can configure ClearPass and Juniper Identity Management Service (JIMS) at the same time. By configuring ClearPass and JIMS at the same time, SRX Series devices can query JIMS for user identification entries, and ClearPass can push device entries to the SRX Series device through the Web API. In releases before Junos OS Release 18.2R1, you are restricted to configure either ClearPass or JIMS.

  [See Understanding How ClearPass and JIMS Works at the Same Time.]

- **Enhancement to NTP authentication method (SRX300, SRX320, SRX340, SRX345, and SRX550M)**—Starting in Junos OS Release 18.2R1, Junos OS supports NTP authentication for both SHA-1 and SHA2-256, in addition to the existing NTP authentication method, MD5. You can now choose from among MD5, SHA-1, and SHA2-256 for synchronizing the clocks of Juniper Network routers, switches, and other security devices on the Internet. Using SHA-1 instead of MD5 improves the security of devices with very little impact to timing, while using SHA2-256 provides an increase in security over SHA-1.

  > **NOTE:** By default, network time synchronization is unauthenticated.

  To implement authentication, use **set authentication-key *<key_number>* type** at the **[edit system ntp]** hierarchy level.

  - To enable SHA-1 authentication, use **set authentication key *<key_number>* type sha1 value *<password>*** at the **[edit system ntp]** hierarchy level.
  - To enable SHA2-256 authentication, use **set authentication key *<key_number>* type sha256 value *<password>*** at the **[edit system ntp]** hierarchy level.

  [See authentication-key and Configuring NTP Authentication Keys.]

Flow and Processing

- **Reverse Route with Packet Mode (SRX Series)**—Starting from Junos OS Release 18.2R1, the reverse route using virtual router is supported with the new CLI command **set security flow advanced-options reverse-route-packet-mode-vr**. While processing the traffic from the server to the client, if the route of the traffic is changed, the traffic is rerouted using the virtual router from the packet incoming interface or filter-based forwarding.

  [See Understanding Reverse Route Packet Mode Virtual Router.]

IDP

- **Flexible grouping of IDP signatures for policies and profiles (SRX Series)**—Starting with Junos OS Release 18.2R1, IDP signature updates support four new tags for creating more sophisticated dynamic groups in addition to the existing seven tags. The signature database is one of the major components of intrusion detection and prevention (IDP). It contains definitions of different objects—such as attack objects, application signatures objects, and service objects—that are used in defining IDP policy rules. Attacks can be grouped by set of tags.

  The additional tags are:

  - CVSS Score (for example, All signatures above 8.0)
  - Age (for example, Older than <x> years)
  - File Type (for example, MPEG, MP4, PPT, and *.doc)
  - Vulnerability Type (for example, buffer overflow, injection, use after free, XSS, and RCE)

  The Product and Vendor tags are already supported under existing filter products. The CLI interface for configuring these tags is now been made more user friendly with possible completions being available for configuration.

  - Vendor (for example, Microsoft, Apple, Red Hat, Google, Juniper, Cisco, and Oracle)
  - Product (for example, Office, Database, Firefox, Chrome, Flash, DirectX, Java, and Kerberos)

[See IDP Policy Rules and IDP Rule Bases.]

Interfaces and Chassis

- **100G Interfaces Support (SRX4600)**–Starting in Junos OS Release 18.2R1, SRX4600 devices support 4x100G Ethernet mode using QSFP28 transceivers. To enable 100-Gigabit Ethernet on the marked ports, use the **set chassis fpc** command.

  [See SRX4600 Gateway Rate-Selectability.]

J-Web

- **J-Web support for Unified L4/L7 Firewall Policy**—Starting Junos OS Release 18.2R1, J-Web supports unified L4/L7 firewall policy, where in you can configure current AppFW by applying its matching criteria of rules to the policy. Also, there are changes to UTM, IPS, AppID, SSL Proxy, Flow and Service redirect.
- **J-Web support for Logical Systems**—Starting Junos OS Release 18.2R1, J-Web supports Logical Systems in SRX5400, SRX5600, and SRX5800 devices, providing multi-tenant firewalls by logically partitioning a single physical firewall into multiple logical systems with separate networking and security services.
- **J-Web support for Configuring ICAP Redirect and SSL Initiation Profiles**—Starting Junos OS Release 18.2R1, using J-Web you can configure ICAP redirect profile and SSL initation profile, which enables you to decrypt HTTPS traffic and redirect HTTP message to 3rd party on-premise DLP server via ICAP/SICAP channel.
- **J-Web Enhanced Look and Feel**—Starting Junos OS Release 18.2R1, J-Web for SRX5400, SRX5600, and SRX5800 devices will have a new and enhanced look and feel.
- **J-Web Configuration Commit Enhancement**—Starting Junos OS Release 18.2R1, after you commit a new J-Web configuration, you can test the configuration for a time period and confirm the commit or roll back to the previous configuration.
- **J-Web support for Logical Domain Interconnect and Routing Instance**—Starting Junos OS Release 18.2R1, using J-Web you can configure the interconnect between logical interfaces and between the root domain and logical systems. Based on the interconnection, you can configure LT interface unit, peer unit, logical system or VPLS switch, and IP addresses for logical system LT interface.

Logical Systems

- **Enabling or disabling ALGs in logical systems (SRX Series)**—Starting in Junos OS Release 18.2R1, you can enable or disable the configuration of Application Layer Gateways (ALGs) in each logical system individually and view the status of the ALGs for all logical systems or specific logical systems. All 12 data ALGs (DNS, FTP, TFTP, MSRPC, SUNRPC, PPTP, RSH, RTSP, TALK, SQL, IKE, and TWAMP) and four VOIP ALGs (SIP, H.323, MGCP, and SCCP) are supported on logical systems.

  [See Understanding Application Layer Gateway (ALG) in Logical System.]

- **Flow enhancement for interconnect logical system (SRX Series)**—Starting in Junos OS Release 18.2R1, the interconnect logical system routing and scaling are supported. You can interconnect multiple logical systems and multiple VPLS switches to pass the traffic without exiting the device. The logical tunnel

interface point-to-point connection **encapsulation frame-relay, encapsulation ethernet** is introduced to optimize the obtainability of logical systems. The frame relay encapsulation adds data-link connection identifier (DLCI) information to the given frame.

[See SRX Series Logical System Master Administrator Configuration Tasks Overview.]

- **Logical systems support (SRX4100 and SRX4200)**—Starting in Junos OS Release 18.2R1, the logical systems are supported on SRX4100 and SRX4200 devices in addition to the existing support on SRX Series devices such as SRX1500, SRX3400, SRX3600, SRX5400, SRX5600, and SRX5800.

  [See Understanding Logical Systems for SRX Series Services Gateways.]

- **Logging support (SRX Series, vSRX)**—Starting in Junos OS Release 18.2R1, the off-box logging (stream mode) service is virtualized. Hence the off-box logging configuration is supported for each logical system and logs are handled based on these configurations. The **[edit logical-system logical-system-name security log]** command is introduced for virtualized logging support. The stream mode is a set of logging services that includes:

  - Off-box logging (SRX Series)

  [See Understanding Security Logs and Logical Systems.]

- **User firewall enhanced support with logical systems (SRX Series)**—Starting in Junos OS Release 18.2R1, support for user firewall authentication is enhanced using a shared model. In this model, user logical systems share user firewall configuration and authentication entries with the root logical system.

  [See Understanding Integrated User Firewall support in a Logical System.]

- **Logical system support (SRX Series)**—Starting in Junos OS Release 18.2R1, SRX4100 and SRX4200 devices support logical system in both transparent and route mode.

  [See Example: Configuring User Logical Systems Security Profiles.]

NAT

- **Network Address Translation (NAT) support for logical systems (SRX Series)**—Starting in Junos OS Release 18.2R1, SRX Series devices support the NAT functionality for logical systems. NAT is a method for modifying or translating network address information in packet headers. Either source or destination addresses or both in a packet can be translated. NAT can include the translation of port numbers as well as IP addresses.

  [See Understanding Logical System Network Address Translation.]

Routing and Forwarding Options

- **NDP and DAD Proxy Support (SRX Series)**—Starting in Junos OS Release 18.2R1, SRX Series devices support Neighbor Discovery Protocol (NDP) and Duplicate Address Detection (DAD) proxy features at the interface level. The NDP and DAD proxies are required if hosts in the same subnet are restricted from communicating directly with each other and need to use the proxy node to forward the packets between them. This feature is primarily used in scenarios where the proxying node needs to apply access control and intercept the traffic flowing between the hosts.

  [See Configuring Duplicate Address Detection Proxy and Configuring Neighbor Discovery Protocol Proxy.]

Security Policies

- **Support for unified policies (SRX Series and , vSRX instances)**—Starting in Junos OS Release 18.2R1, unified policies are now supported on SRX Series devices and vSRX instances, allowing granular control and enforcement of dynamic Layer 7 applications within the traditional security policy.

  Unified policies are the security policies, where you can use dynamic applications as match conditions along with existing 5-tuple or 6-tuple matching conditions (with user firewall) to detect application changes over time, and allow you to enforce a set of rules for the transit traffic.

  Unified policies allow you to use dynamic application as one of the policy match criteria rule in each application. Application identification (AppID) is applied on the traffic, and the application is identified after several packets are checked.

  Before identifying the final application, the policy cannot be matched precisely. A potential policy list is made available, and the traffic is permitted using the potential policy from the list.

  After the application is identified, the final policy is applied to the session. Policy actions such as permit, deny, reject, or redirect is applied on the traffic as per the policy rules.

  [See Understanding Unified Policies.]

The following features support unified policies:

- **Application Identification (AppID)**—Unified policy leverages the application identity information from the Application Identification (AppID). AppID provides the information such as dynamic application classification, default protocol and port of an application. For any application included in the dependent list of another application, AppID provides this information.

  [See Application Identification.]

- **Application firewall (AppFW)**—Unified policy configuration handles AppFW functionality and simplifies the task of configuring firewall policy to permit or block application traffic from the network.

  If you configure a unified policy with a dynamic application as one of the matching conditions, then the configuration eliminates the additional steps involved in AppFW configuration—that is, configuring a security policy to invoke the application firewall service.

  Starting in Junos OS Release 18.2R1, the Application Firewall (AppFW) functionality is deprecated—rather than immediately removed—to provide backward compatibility and a chance to bring your configuration into compliance with the new configuration.

  The **[edit security application-firewall]** hierarchy level and all configuration options under this hierarchy are deprecated.

  [See Application Firewall.]

- **Application Quality of Service (AppQoS)**—AppQoS functionality is supported when the device is configured with unified policies. You can configure a default AppQoS rule set to manage unified policy conflicts, if multiple security policies match the traffic.

  [See Application Quality of Service.]

- **ICAP service redirect**—Internet Content Adaptation Protocol (ICAP) service redirect functionality is supported when the device is configured with unified policies.

  [See iCAP Service Redirect.]

- **IDP**—Starting with Junos OS Release 18.2R1, with unified policies support, when a security rule has IDP enabled, the name of the actual IDP policy is replaced. This is to simplify IDP policy usage and to provide flexibility to have multiple policies active at the same time.

  All IDP matches will now be handled within the unified policies. As a part of session interest check IDP will enabled if IDP policy is present in any of the matched rules.

  IDP policy is activated in security policies, by permitting the IDP policy within the application services using the **set security policies from-zone zone-name to-zone zone-name policy policy-name then permit application-services idp-policy idp-policy-name** command.

  Since IDP policy name is directly use in the security policy rule, the **[edit security idp active-policy policy-name]** statement is deprecated.

  [See IDP Policies Overview.]

- **SSL proxy**—SSL proxy functionality is supported when the device is configured with unified policies. You can configure a default SSL proxy profile to manage unified policy conflicts, if multiple security policies match the traffic.

  [See SSL Proxy.]

- **UTM**—A new dynamic application policy match condition is added to SRX Series devices, allowing an administrator to more effectively control the behavior of Layer 7 applications. To accommodate Layer 7 application-based policies in UTM, the **[edit security utm default-configuration]** command is introduced. If any parameter in a specific UTM feature profile configuration is not configured, then the corresponding parameter from the UTM default configuration is applied.

  Additionally, during the initial policy lookup phase which occurs prior to a dynamic application being identified, if there are multiple policies present in the potential policy list which contains different UTM profiles, the SRX Series device applies the default UTM profile until a more explicit match has occurred.

  [See Understanding Unified Policies [Unified Threat Management (UTM)].]

- **Juniper Sky ATP support within unified policy (SRX Series)**— Juniper Sky ATP is supported for unified policies. The **set services security-intelligence default-policy** and **set services advanced-anti-malware default-policy** commands are introduced to create default settings for both policy types. During the initial policy lookup phase, which occurs prior to a dynamic application being identified, if there are multiple policies present in the potential policy list, which contain different security intelligence or anti-malware policies, the SRX Series device applies the default policy until a more explicit match has occurred.

  [See the Juniper Sky ATP Administration Guide.]

User Interface and Configuration

- **Support for displaying ephemeral configuration data with filtering (SRX Series)**—Starting in Junos OS Release 18.2R1, the **show ephemeral-configuration** command enables you to specify the scope of the configuration data to display. To filter the displayed configuration data, append the statement path of the requested hierarchy to the command.

  [See Displaying Ephemeral Configuration Data in the Junos OS CLI.]

UTM

- **Antispam supports IPv6 address [SRX Series]** —Starting in Junos OS Release 18.2R1, the antispam feature supports IPv6 traffic.

  [See Antispam Filtering.]

VPN

- **Configuring forwarding class on IPsec VPNs (SRX Series, vSRX)**—Starting with Junos OS Release 18.2R1, forwarding classes configured on an SRX Series device can be mapped to IPsec security associations

(SAs). Multiple IPsec SAs are negotiated on the same IKE SA with a peer device, one SA per forwarding class configured in IPsec.

A unique IPsec SA is negotiated with the VPN peer for each forwarding class. By mapping the forwarding class to the IPsec SA, all the packets with a certain class-of-service (CoS) value will get quality-of-service (QoS) treatment between the peer devices thus avoiding packet drop due to the anti-replay window. This feature provides QoS for IPsec when peer devices allow for multiple SA negotiation.

[See Understanding CoS-Based IPsec VPNs with Multiple IPsec SAs.]

- **Public key infrastructure (PKI) proxy support (SRX Series)**—Starting in Junos OS Release 18.2R1, PKI supports Hypertext Transfer Protocol (HTTP) Web proxy. HTTP Web proxy acts as an intermediary between the client and the server, but neither the server nor the client can detect its presence. You can add Web proxy support to the SRX Series devices to configure systemwide HTTP connections to the egress traffic to ensure secure communication with the certificate authority (CA) server.

  [See Understanding Certificate Authority Profiles.]

SEE ALSO

Changes in Behavior and Syntax

This section lists the changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands from Junos OS Release 18.2R3 for the SRX Series.

Juniper Sky ATP

- **Dynamic address entries on SRX Series devices in chassis cluster mode**—Starting in Junos OS Release 18.2R3, for SRX Series devices in chassis cluster mode, the dynamic address entry list is retained on the device even after the device is rebooted following a loss of connection to Juniper Sky Advanced Threat Prevention (ATP).

Network Management and Monitoring

- **NSD Restart Failure Alarm (SRX Series)**—Starting in Junos OS Release 18.2R3, a system alarm is triggered when the Network Security Process (NSD) is unable to restart due to the failure of one or more NSD subcomponents. The alarm logs about the NSD are saved in the messages log. The alarm is automatically cleared when NSD restarts successfully.

  The **show chassis alarms** and **show system alarms** commands are updated to display the following output when NSD is unable to restart - **NSD fails to restart because subcomponents fail**.

  [See Alarm Overview.]

- Starting with Junos OS Release 18.2R1, the following commands under the **[edit security utm feature-profile]** hierarchy level are deprecated:

  - set web-filtering type
  - set web-filtering url-blacklist
  - set web-filtering url-whitelist
  - set web-filtering http-persist
  - set web-filtering http-reassemble
  - set web-filtering traceoptions
  - set web-filtering juniper-enhanced cache
  - set web-filtering juniper-enhanced reputation
  - set web-filtering juniper-enhanced query-type
  - set anti-virus mime-whitelist
  - set anti-virus url-whitelist
  - set anti-virus type
  - set anti-virus traceoptions
  - set anti-virus sophos-engine
  - set anti-spam address-blacklist
  - set anti-spam address-whitelist

- set anti-spam traceoptions
- set content-filtering traceoptions

[See feature-profile.]

Security

- Starting with Junos OS Release 18.2R1, the following commands under the **[edit security utm feature-profile]** hierarchy level are deprecated:
  - set web-filtering type
  - set web-filtering url-blacklist
  - set web-filtering url-whitelist
  - set web-filtering http-persist
  - set web-filtering http-reassemble
  - set web-filtering traceoptions
  - set web-filtering juniper-enhanced cache
  - set web-filtering juniper-enhanced reputation
  - set web-filtering juniper-enhanced query-type
  - set anti-virus mime-whitelist
  - set anti-virus url-whitelist
  - set anti-virus type
  - set anti-virus traceoptions
  - set anti-virus sophos-engine
  - set anti-spam address-blacklist
  - set anti-spam address-whitelist
  - set anti-spam traceoptions
  - set content-filtering traceoptions

[See feature-profile.]

VPNs

- **Certificate revocation list (SRX Series)**—Local certificates are being validated against certificate revocation list (CRL) even when CRL check is disabled. Starting in Junos OS Release 18.2R3, this can be stopped by disabling the CRL check through the Public Key Infrastructure (PKI) configuration. When CRL check is disabled, PKI will not validate local certificate against CRL.

  [See revocation-check (Security PKI) and Understanding Online Certificate Status Protocol and Certificate Revocation Lists.]

SEE ALSO

New and Changed Features | 342

Known Behavior | 357

Known Issues | 360

Resolved Issues | 363

Documentation Updates | 380

Migration, Upgrade, and Downgrade Instructions | 380

Product Compatibility | 381

Known Behavior

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.2R3 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.

Chassis Clustering

- On all SRX branch devices, 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 sends 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 drops. PR1344173

Flow-Based and Packet-Based Processing

- On SRX4600 devices, the USB disk is not available for Junos OS. However, the USB disk is available for the host OS (Linux) with full access. The USB disk is still used in the booting process (install and recovery functions). PR1283618
- When a USB device is under initialization, removing the USB device leads to USB crash. PR1332360
- On SRX1500 devices with AppFW configured, the expected HTTP CPS is 60,000, which is a 14 percent drop (the expected value is 70,000). PR1339131
- The memory usage of useridd increases when the configuration is exchanged between user firewall for active directory and JIMS. PR1383751
- Z mode is not supported when using the dedicated fabric link. PR1397267
- On an SRX high-end platform, if a huge packet is bigger than 9271 in size and is fragmented, when it hits the ALG gate, the defragmented huge packet needs to be forwarded to another SPU. There is a size limitation when forwarding a packet between SPUs, which would cause the huge packet to be dropped in such situations. PR1426644
- If incoming tunnel and egress tunnel are anchored on different SPUs, PMI flaps this x2 traffic to regular flow path for encryption on the egress tunnel because PMI does not have the PIC forwarding functionality. PR1432915

J-Web

- On SRX Series devices, DHCP relay configuration on the **Configure > Services > DHCP > DHCP Relay** page is removed from the J-Web interface 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 on the **Monitor** page is removed. You can view the same bindings in the CLI using the **show dhcp client binding** command. PR1205915
- On SRX Series devices, adding 2000 global addresses at a time to the SSL proxy profile exempted addresses can cause the Web page to become unresponsive. PR1278087
- On SRX Series devices, you cannot view the custom log files created for event logging using J-Web interface. PR1280857

- Uploading a certificate using the **Browse** button stores the certificate in the device at the **/jail/var/tmp/uploads/** location, which will be deleted upon executing the **request system storage cleanup** command. PR1312529
- The values of address and address-range are not displayed in the inline address-set creation pop-up window of the Juniper Identity Management Service (JIMS) server. PR1312900
- PPPoE interface pp0 will not be manageable through J-Web **Interfaces->Port** page. PR1316328
- The **Dynamic-Application** configuration page in the J-Web interface did not properly display application-signatures when searching by **category**. PR1344165
- Forming an HA from J-Web by using the HA cluster wizard is not supported from Junos OS Release 12.1X47 onward for SRX5400 only. PR1372518

Routing Protocols

- A new CLI command is required to prevent traffic loss during a disaster recovery failover scenario. PR1352589

User Interface and Configuration

- On SRX Series devices setups, committing a configuration with a considerable number of logical system configuration can take a little more time than usual. This issue occurs because backing up previous configurations might take a little longer to finish. PR1339862

VPNs

- When multiple traffic selectors are configured on a particular VPN, the iked process checks for a maximum of one DPD probe that is sent to the peer for the configured DPD interval. The DPD probe is sent to the peer if traffic flows over even one of the tunnels for the given VPN object. PR1366585
- The iked process can handle 1000 DPD packets per second. If HA link encryption is enabled, the iked process can handle 500 DPD packets per second. DPD packets include both DPD probes sent from the device and DPD probes received from the peer. PR1380971
- Use the file created in the **set security ike traceoptions file** location to check the logs. PR1381328
- In the output of the **show security ipsec inactive-tunnels** command, **Tunnel Down Reason** is not displayed as this functionality is not supported in Junos OS Release 18.2R2 and later. PR1383329

- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, occasionally, if an IKE or IPsec configuration (under groups hierarchy) is changed for one IKE gateway, the tunnel might be cleared for unrelated IKE or IPsec gateways. PR1405840
- The iked process does not handle cases and core files might be generated when a remote gateway address is configured as an IPv6 address while the local interface where the tunnel is anchored has an IPv4 address. PR1416081

SEE ALSO

New and Changed Features | **342**

Changes in Behavior and Syntax | **354**

Known Issues | **360**

Resolved Issues | **363**

Documentation Updates | **380**

Migration, Upgrade, and Downgrade Instructions | **380**

Product Compatibility | **381**

Known Issues

This section lists the known issues in hardware and software in Junos OS Release 18.2R3 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 Security

- If automatic application-identification download is configured with a start-time specified, the automatic download stops when the time has progressed to the next year and a reboot is done before the start-time is reached that year. PR1436265

Application Identification

- IDP installation fails on one node due to application identification process stuck. PR1336145

Flow-Based and Packet-Based Processing

- On SRX4600 platform, the output of the **show route forwarding-table** command displays the **next-hop** IP address twice if the **next-hop** is the st0 interface. The routing functionality is not impacted. PR1290725
- The device sends incorrect rejection code when the destination device is not reachable. PR1371115
- An SRX Series device receives a dynamic update from the JIMS server when JIMS detects that a user has been disabled in the Active Directory, and another dynamic update from JIMS if that user is subsequently reenabled. This update retains the domain and username in the table, but does not include any groups associated with that user. PR1380771
- With stress TCP traffic, sessions that have been invalid for more than 48 hours expire. PR1383139
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, when multiple cores occur in quick succession, the coldsync monitored status is displayed and it is not removed after completing the coldsync. You need to reboot the affected node to recover from this issue. PR1403000
- On all SRX Series platforms, in chassis cluster with Z mode traffic and local (non-reth) interfaces 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 might get dropped due to reroute failure. PR1410233
- When a GRE tunnel (GRE over IPsec tunnel) or IPsec tunnel is used on an SRX Series device, the MTU of the tunnel interface is calculated incorrectly (24 bytes less than the expected value). PR1426607
- On SRX5000 Series devices with an SPC3 card, traffic might not go through when BA classifiers for DSCP are configured with rewrite. PR1428153
- On SRX5000 Series devices with an SPC3 card, sometimes IKE SA is not seen on the device when st0 binding on VPN configuration object is changed from one interface to another (for example, st0.x to st0.y) PR1441411

J-Web

- On SRX Series platforms, the root password configured at first J-Web access (**Skip to J-Web** feature) does not work if the password length is shorter than eight characters. PR1371353

Platform and Infrastructure

- On SRX5600 and SRX5800 devices in chassis cluster, when a second Routing Engine is installed to enable dual control links, the **show chassis hardware** command might display the same serial number for both the secondary Routing Engines on both the nodes. PR1321502
- On SRX5000 platforms (include SRX5400, SRX5600,SRX5800), the 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

VPNs

- When an SRX Series device acts as an initiator behind the NAT, disabling NAT on the router in between causes an immediate new negotiation failure because of an attempt to disable NAT using port 4500. The next attempt succeeds by using port 500. Disabling NAT and bringing down all the existing tunnels and reestablishing the tunnels with port 500 is the expected behavior. PR1273213
- On SRX Series devices, if multiple traffic selectors are configured for a peer with Internet Key Exchange version 2 (IKEv2) reauthentication, only one traffic selector is rekeyed at the time of IKEv2 reauthentication. The VPN tunnels of the remaining traffic selectors are cleared without immediate rekeying. A new negotiation of these traffic selectors is triggered through other mechanisms—for example, by traffic or by a peer. PR1287168
- When using the operational mode **request security ike debug-enable** for IKE debugging after having used IKE traceoptions with a file name specified in the configuration, the debugs are still being written to the same file name. PR1381328
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, when a large number of IPsec tunnels are established, a few tunnels might fail during rekey negotiation if the device initiates the rekey. PR1389607
- VPN does not recover on the high-end standalone SRX Series devices when CLI operation **restart ipsec-key-management** is done. PR1390831
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, occasionally, if the IKE or IPsec configuration (under groups hierarchy) is changed for one IKE gateway, the tunnel might clear for an unrelated IKE or IPsec gateway. PR1405840
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, a new behavior has been introduced that differs from the behavior on the older SPC2 card. The SRX device with AutoVPN configuration can

now accept multiple IPsec tunnels from a peer device (with the same source IP address and port number) using different IKE-IDs. PR1407356

- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, if an existing IKE gateway configuration is changed from AutoVPN to Site-to-Site VPN, the IKE negotiation behavior remains in **responder-only** mode. PR1413619
- On SRX5400, SRX5600, and SRX5800 devices, during in-service software upgrade (ISSU), the IPsec tunnels flap, disrupting traffic. The IPsec tunnels recover automatically after the ISSU process is completed. PR1416334
- On SRX5000 Series devices with an SPC3 or SPC2 card, when a duplicate user (same user with different IP address but same IKE-ID) logs in, in some cases, old IKE SA entries are not deleted immediately. PR1423821
- IKE SA entries are not displayed in CLI output after a cluster node fails over when tunnels are established in aggressive mode. PR1424077

SEE ALSO

New and Changed Features | **342**

Changes in Behavior and Syntax | **354**

Known Behavior | **357**

Resolved Issues | **363**

Documentation Updates | **380**

Migration, Upgrade, and Downgrade Instructions | **380**

Product Compatibility | **381**

Resolved Issues

IN THIS SECTION

- Resolved Issues: 18.2R3 | **364**
- Resolved Issues: 18.2R2 | **370**
- Resolved Issues: 18.2R1 | **377**

This section lists the issues fixed in Junos OS Release 18.2R3 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.

Resolved Issues: 18.2R3

Application Identification

- Application identification classification logic has been improved for NetBIOS and RPC. PR1357093

Application Layer Gateways (ALGs)

- DNS requests with additional EDNS records might be dropped by the DNS ALG. PR1379433
- On all SRX Series devices, SIP/FTP ALG does not work when SIP traffic with source NAT goes through the SRX Series device. PR1398377
- The H.323 protocol voice packets might be dropped on SRX Series devices. PR1400630
- When both ALG and **rst-invalidate-session** are enabled, the TCP reset packet is dropped by SRX Series devices. PR1430685

Application Security

- On all SRX Series device with Security Intelligence (SecIntel), the Black/White list file open might fail if the file pointer is null, which might cause the ipfd process to stop. PR1436455
- On all SRX Series devices with advanced anti-malware service used, due to a rare issue in file system handling in data plane, the flowd/srxpfe process might crash. PR1437270

Chassis Clustering

- The SNMP trap sends wrong information with manual failover. PR1378903
- GTPv2 modify bearer request packet that does not contain F-TEID IE in bearer context is dropped during GTP inspection. PR1399658
- Traffic with domain name address might fail for 3-5 minutes after RG0 failover on SRX Series devices. PR1401925
- The flowd process stops while updating or deleting a GTP tunnel. PR1404317
- Mixed-mode (SPC3 card coexisting with SPC2 cards) HA IP monitoring fails on secondary node with the following error message: **secondary arp entry not found**. PR1407056
- SRX Series devices might be potentially overwritten with an incorrect buffer address when detailed logging is configured under GTPv2 profile. PR1413718
- PIM neighbors might not come up on a chassis cluster on SRX Series devices. PR1425884

- Starting in Junos OS Release 18.4, a maximum of six PDN connects can be contained in PDP context response. Otherwise, the response is dropped. PR1422877
- RG0 failover sometimes causes FPC offline/present status. PR1428312

Flow-Based and Packet-Based Processing

- Removal of RC4 from SSL-FP. PR1302789
- SRX1500 device 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
- Control traffic loss might be seen on SRX4600 platform. PR1357591
- When activating **security flow traceoptions**, the unfiltered traffic is captured. PR1367124
- SRX1500 device continues alarm on **Fan Tray 0 Fan 0 Spinning Degraded**. PR1367334
- On SRX1500 device, the activity LED (right LED) for 1-Gigabit Ethernet or 10-Gigabit Ethernet port is not on although traffic is passing through that interface. PR1380928
- Password recovery menu does not show up on SRX Series devices. PR1381653
- Large file downloads slow down for many seconds. PR1386122
- Traffic might be processed by the VRRP backup when multiple VRRP groups are configured. PR1386292
- The default configuration of SRX300 line of devices is changed. PR1393683
- Switching interface mode between **family ethernet-switching** and **family inet/inet6** might cause traffic loss. PR1394850
- Performance drops are seen in SRX345 and SRX340 platforms for IDP C2S policy. PR1395592
- These messages are seen: **kernel: tcp_timer_keep:Local(0x80000004:54652) Foreign(0x80000004:33160)**. PR1396584
- On SRX4600 platform, the 40 Gigabit Ethernet interface might flap continuously because of a MAC local fault. PR1397012
- 40-Gigabit Ethernet or 100-Gigabit Ethernet ports may take a long time (about 30 seconds) to link up on SRX4600 platform. PR1397210
- The pkid process might stop after RG0 failover. PR1379348
- On SRX Series devices, the connection to JIMS fluctuates, resulting in failover. PR1398140
- On SRX4600 and SRX5000 Series devices, BGP packets might be dropped when CPU usage is high. PR1398407
- SRX Series devices do not strip VLAN added by native **vlan id** command options. PR1397443
- The next-hop IP address is not displayed in the routing table in the J-Web interface. PR1398650
- VLAN push might not work on SRX1500. PR1398877
- Increase DAG feed scale number to 256 from 63. PR1399314

- The authd process might stop when the **show network-access requests pending** command is issued while the authd process is restarting. PR1401249
- SRX Series device cannot obtain IPv6 address through DHCPv6 when using PPPOE interface with logical-unit-number greater than zero. PR1402066
- Unable to access SRX Series devices if messages **kern.maxfiles limit exceeded by uid 65534, please see tuning(7)** are seen. PR1402242
- CPU is hitting 100 percent with fragmented traffic. PR1402471
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, when PowerMode IPsec is enabled, the **show security flow statistics** and **show security flow session tunnel summary** command do not count or display the number of packets processed within PowerMode IPsec because these packets do not go through regular flow path. PR1403037
- Downloads might stall or fail completely when utilizing services that are reliant on TCP proxy. PR1403412
- Transit UDP 500/4500 traffic might not pass across SRX5000 Series devices when using an SPC3 or SPC2 card. PR1403517
- Split brain condition is experienced if the an SPC3 or SPC2 card goes offline in the primary node. PR1403872
- Fail to match permit rule in application firewall ruleset. PR1404161
- Configuring using the CLI editor in the J-Web generates an mgd core file. PR1404946
- The flowd process stops and all cards are brought offline. PR1406210
- The RG1 failover does not happen immediately when the SPC3 card crashes. PR1407064
- The flowd process might crash if **enable-session-cache** command is configured under the SSL termination profile. PR1407330
- IDP signature update fails at RG0 primary node. PR1407603
- On SRX1500 platform, traffic is blocked on all interfaces after configuring **interface-mac-limit** on one interface. PR1409018
- Memory leak occurs if AAMW is enabled. PR1409606
- While PMI is ON, IPsec encrypted statistics on the Routing Engine **show security ipsec statistics** is not working anymore for fragment packets. PR1411486
- PEM 0 or PEM 1 I2C failure major alarm might be set and cleared for multiple times. PR1413758
- HA packets might be dropped on SRX5000 Series devices with an IOC3 or IOC2 card. PR1414460
- Any traffic originated from the device itself might be dropped in the IPsec tunnel. PR1414509
- The input and output bytes or bps statistic values might not be identical for the same size of packets. PR1415117
- Traffic is dropped if SOF is enabled in chassis cluster active/active mode PR1415761

- Juniper Sky ATP does not escape the \ inside the username before the metedata is sent to cloud. PR1416093
- The flowd process stops on SRX5000 or SRX4000 Series devices when large-size packets go through IPsec tunnel with the post-fragment check. PR1417219
- Traffic logging shows **service-name junos-dhcp-server** for UDP destination port 68. PR1417423
- Traffic might be lost on SRX Series devices if IPsec session affinity is configured with **ipsec-performance-acceleration**. PR1418135
- SSL proxy does not correctly warn users about unsupported certificates. PR1419485
- FRU model number is not displayed. PR1422185
- The **show security flow session session-identifier < sessID>** command does not work if the session ID is bigger than 10M on SRX4600 platform. PR1423818
- Partial traffic might get dropped on an existing LAG. PR1423989
- Memory leaks might occur on the jsqlsyncd process on SRX chassis clusters. PR1424884
- Alarms due to high temperature when operating with expected temperatures. PR1425807
- The IPsec traffic going through SRX5000 Series device with SPC2 cards installed causes high SPU CPU utilization. PR1427912
- Uneven distribution of CPU with high PPS on device. PR1430721
- The flowd process might stop on SRX5000 Series devices. PR1430804
- SRX550M running Junos OS Release 18.4R1 shows PEM 1 output failure message where as with Junos OS Release 15.1X49 or Junos OS Release 18.1R3.3 it does not show any alarms. PR1433577
- SPMC version mismatch errors after Junos OS is installed using USB method. PR1437065
- Performance improvements were made to Screens which benefit multi-socket systems like the SRX4200 and SRX4600 devices, and SPC3's. PR1440677
- On SRX5400, SRX5600, SRX5800 Series platforms acting as a middle device between Internet Key Exchange (IKE) peers, it is not able to establish more than one Encapsulating Security Payload (ESP) session between two IPv6 IKE peer if the IKE ALG is enabled on the middle SRX device. PR1435687

Interfaces and Chassis

- Both nodes in the SRX cluster went into db mode after downgrading to Junos OS Release 18.1 when the **vlan-tagging** configured on reth interfaces, but vlan-id is not configured PR1407295

Intrusion Detection and Prevention (IDP)

- IDP might stop with the custom IDP signature. PR1390205
- Unable to configure **dynamic-attack-group**. PR1418754

J-Web

- The **Dynamic-Application** configuration page within J-Web does not properly display application-signatures when searching by **category**. PR1344165
- In the J-Web dashboard, the **Security Resources** widget does not display absolute values. PR1372826
- J-Web now supports defining SSL proxy and redirect (block page) profiles when a policy contains dynamic applications. PR1376117
- Special character used in the pre-shared-key is removed silently after a commit operation on J-Web. PR1399363
- The httpd-gk process stops leading to dynamic VPN failures and high Routing Engine CPU utilization up to 100 percent. PR1414642
- J-Web configuration change for address set using **Search** function results in commit error. PR1426321
- J-Web shows incorrect port-mode under **Configure>Interfaces>Link Aggregation**. PR1430414
- IRB interface is not available in zone option of J-Web. PR1431428

Multiprotocol Label Switching (MPLS)

- The rpd process might restarts unexpectedly when **no-cspf** and lo0 is not included under RSVP. PR1366575

Network Address Translation (NAT)

- SRX-SPC3 mix mode NAT SPC3 core file at **../sysdeps/unix/sysv/linux/raise.c:55**. PR1403583
- The nsd process stops and causes the Web filter to stop working. PR1406248

Network Management and Monitoring

- The **set system no-redirects** setting does not take effect for reth interface. PR894194
- The chassisd might stop and restart after the AGENTX session timeout between master (snmpd) and sub-agent. PR1396967

Platform and Infrastructure

- High httpd utilization after reboot failover. PR1352133
- Log in class with allowed-days and specific access-start/access-end does not work as expected. PR1389633
- Memory leak might occur on the data plane during composite next-hop installation failure. PR1391074
- GW lcores and srxpfe cores files 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
- Complete device outage might be seen when SPU VM core happens. PR1417252

- Some packages might be omitted during upgrade from legacy with packages. PR1417321
- Flowd process might stop on SRX Series devices. PR1417658
- Routing Engine CPU utilization is high and eventd consumes a lot of resources. PR1418444
- On SRX4600 device, commit failed while configuring 2047 VLAN IDs on reth interface. PR1420685
- The interface flaps when LACP is configured on the reth interfaces. PR1435955

Routing Policy and Firewall Filters

- The **show security flow session** command now fully supports the dynamic-application construct. PR1387449
- Memory leak in nsd prevents the configuration from taking effect after it is committed. PR1414319
- The output of the **show security firewall-authentication jims statistics** command displays the statistics of both the primary JIMS server and secondary JIMS server. PR1415987
- The flowd process stops while deleting policies from Junos Space. PR1419704
- A commit warning is displayed when a traditional policy is placed below a unified policy. PR1420471
- One new alarm is created: **NSD fails to restart because subcomponents fail**. PR1422738

Unified Threat Management (UTM)

- Whitelist/Blacklist does not work for HTTPS traffic going through Web proxy. PR1401996
- On SRX Series devices, when configuring Enhanced Web Filtering on the CLI, the autocomplete function does not properly handle or suggest custom categories. PR1406512
- The device might not look up blacklist first in local Web Filtering environment. PR1417330
- UTM Web filtering status shows down when using Hostname [routing-instance synchronization failure]. PR1421398
- When using unified policies, the base-filter for certain UTM profiles might not be applied correctly. PR1424633
- Behavioral improvements were made to SSL-Proxy's url-category whitelisting functionality. PR1426189

VPNs

- SPC3 IKE SA detail output does not show proper traffic statistics. PR1371638
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, the **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
- VPN does not recover on the high end standalone SRX Series devices when CLI operation **restart ipsec-key-management** is done. PR1400712

- Syslog is not generated when IKE gateway rejects duplicate IKE ID connection. PR1404985
- Idle IPsec VPN tunnels without traffic and with ongoing DPD probes might be affected during RG0 failover. PR1405515
- 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
- On SRX5400, SRX5600, SRX5800 devices with SPC3, when SRX is configured in IKEv1 and NAT traversal is active, after a successful IPsec rekey IPsec tunnel index may change. In such a scenario, there might be some traffic loss for a few seconds. PR1409855
- Traffic drops on peer due to bad SPI after first reauthentication. PR1412316
- On SRX5400, SRX5600, SRX5800 devices with an SPC3 card, when the device is configured to initiate IKEv2 reauthentication when NAT traversal is active, occasionally reauthentication might fail. PR1414193
- The flowd/srxpfe process might crash when traffic selector is used for IPsec VPN. PR1418984
- The **show security ike sa detail** command shows incorrect value in **IPSec security associations** coloumn. PR1423249
- Once VPN IPsec with NATT (NAT in the middle of IPSEC peers) is in place, the SRX Series devices performance is slow. PR1424937
- SRX Series devices should send IKE delete notification to peer when traffic selector configuration is changed for a specific AutoVPN. PR1426714
- Kmd process stops and generates a core file after running the CLI command **show security ipsec traffic-selector <>**. PR1428029
- VPN overhead calculation is going wrong on SPC3 due to using wrong spu-id API. Fixed this issue by calling common API for SPC2 and SPC3 to get SPU-id without core-id. PR1435700

Resolved Issues: 18.2R2

Application Layer Gateways (ALGs)

- On SRX320, SRX340, SRX340, and SRX550 devices, the RPD process stops when you configure the **auto-bandwidth** option under the label-switched path (LSP) in the multiprotocol label switching (MPLS). PR1331164
- When using the IPsec ALG, the IPsec tunnel payload is dropped after the IKE or IPsec tunnel reestablishment due to a session conflict. PR1372232
- If the SIP ALG is disabled, the SIP active sessions are affected. PR1373420
- DNS requests with EDNS option might be dropped by the DNS ALG. PR1379433

- SUN RPC data traffic for previously established ALG sessions might be dropped because it matches the gate which contains old interface information. PR1387895
- On SRX5400, SRX5600, and SRX5600 devices, flowd process might generate core files while sending cross tenant ALG traffic. As a workaround, avoid the cross tenant ALG traffic or disable ALG type which has cross tenant traffic. PR1388658

Chassis Clustering

- On SRX340 and SRX345 devices, half-duplex mode is not supported because BCM53426 does not support half-duplex mode. BCM5342X SoC port configurations, BCM53426 does not have QSGMII interface. Only the QSGMII port supports halfduplex mode. PR1149904
- On SRX550M device, the SFP transceiver does not work after the chassis reboot. PR1347874
- On SRX4600 device with chassis cluster enabled, when a failover occurs the dedicated fabric link is down. PR1365969
- The device in chassis cluster might be unresponsive if IP monitoring is enabled. PR1366958
- On SRX Series devices in chassis cluster, minor **Potential slow peers are: FWDD0 XDPC1 XDPC8 FWDD1** alarm is observed which can be ignored. PR1371222
- Multiple flowd process files are seen on node 1 after an RG0 failover. PR1372761
- Traffic loss occurs when the primary node is rebooting. PR1372862
- On SRX Series devices in chassis cluster, if reroute occurs on the IPv4 wings of a NAT64 or NAT46 session, the active node will send RTO message to the backup session to update the rerouted interface. PR1379305
- On SRX4600 device in chassis cluster, the Flexible PIC Concentrators (FPCs) goes offline if the chassis clusters IDs are more than 10. PR1390202

Class of Service (CoS)

- When the **host-outbound-traffic** command is configured in the class of service (CoS), the device stops working when a corrupted packet is arrived on the Packet Forwarding Engine. PR1359767

Command Line Interface

- The following CLI command outputs are not displayed correctly: **show usp memory segment shm data module** and **show jsf shm module**. PR1387711

Dynamic Host Configuration Protocol

- SRX300, SRX320, SRX340, and SRX345 devices with LTE mPIM do not forward the DHCP relay packets over the LTE. PR1357137

Flow-based and Packet-based Processing

- The security logs for unified policies are improved to reflect the reason for a denied or rejected session. PR1338310
- On SRX Series devices, the **session-init** and **session-close** are logged for unified policies. PR1338319
- The IPsec replay error for Z-mode traffic is observed. PR1349724
- Multicast routes are not seen after setting the maximum transmission unit (MTU) size to 1,300. PR1349996
- On SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 devices in chassis cluster, when CoS is configured on a interface, LACP communication stops due to failure of the fabric port and the connections between the SRX device and other devices breaks. PR1350731
- When the routing instance is configured, the UTM antispam does not send the DNS query. PR1352906
- The IPsec VPN traffic might be dropped on the pass-through device after an IKE rekey operation. PR1353779
- The PIM register message from source first-hop router (FHR) suddenly stops appearing. PR1356241
- When the output interface configured in X2 mirror-filter is down, the flowd process might stop. PR1357347
- On SRX4200 and SRX4600 devices, during reboot or power on the device, the control traffic loss is observed. PR1357591
- IDP inline tap mode is not supported and configuration for SPC3 must be disabled. PR1359591
- On the secondary control plane, a multicast session leak is observed for the PIM register. PR1360373
- On SRX Series devices, if you disable one of the four reth interfaces, the traffic flow stops. PR1360399
- On SRX Series devices, in an SSL proxy scenario, if the TLS packets contain application layer protocol negotiation (ALPN), then the ALPN extension is removed by the SSL proxy, resulting in the negotiation failure of the application layer protocol. PR1360820

- The default version of Application identification (AppID) signature pack and protocol bundle are updated. PR1362367
- On SRX550M device, the traffic might be duplicated and forwarded to the wrong interface. PR1362514
- This release includes support for service-specific ASC configuration, which allows the ASC to be enabled or disabled on a per-service basis. With the advent of Unified Policies, two services are introduced to the ASC: Security Services, and Miscellaneous Services. Security Services are responsible for policy-lookup behavior, while miscellaneous services are responsible for non-policy related items, such as APBR. Starting with 18.2R1, by default, the ASC will be disabled for security-services and enabled for miscellaneous-services. This has the possibility to impact existing legacy AppFW functionality post-upgrade to 18.2R1 as existing cache-entries will be ignored during policy-lookups. PR1363501
- SNMP MIB walk does not work when screens are applied to more than 14 security zones. PR1364210
- On SRX Series devices, the Application identification (AppID) is supported for HTTP, SMTPS, POP3S and IMAPS protocols. PR1365810
- On SRX5400, SRX5600 and SRX5800 devices with chassis cluster, in-service software upgrade (ISSU) might struck at **IPID data synn** state. PR1366077
- A flowd core file is generated after an RG0 failover. PR1366122
- On SRX Series devices, when Application Quality of Experience (AppQoE) is enabled and the traffic starts flowing, the flowd process might stop. PR1367599
- On SRX1500 device with Junos OS 15.1X49-D140, the srxpfe process might not work. PR1370900
- The SPC3 core file size is larger than the SPC1 and SPC2 core dump files. PR1371447
- In chassis cluster mode with the IPsec tunnel configured, packet loss is observed when the clear text packets are processed. PR1373161
- The SPC3 card improved the unified policies performance. PR1374231
- On SRX Series devices, the Security Log Event Details window size is increased to display all the relevant information about the event. PR1373357
- On SRX Series devices working in a PIM sparse mode, the network located between a first hop router (FHR) and a rendezvous point (RP), if a PIM control session is created through the PIM register stop message, only the next one PIM register message can be forwarded, after this message, the subsequent PIM register message (also matches the above PIM control session) is wrongly dropped. PR1378295
- When the data path debug capture is stopped, incorrect error message is displayed. PR1381703
- The SPC3 might be installed on any slot except slot 0, slot 1, and slot 11. PR1378178
- On SRX5600 with chassis cluster, if the respmod is enabled for ICAP, the connection with the ICAP server might reset automatically. PR1382376
- On SRX300, SRX320, SRX340, SRX345, SRX550-M devices, during the path MTU discovery, CE does not receive the message **frag needed and DF set**. PR1389428

- On SRX4600, SRX5400, SRX5600, and SRX5800 devices using the SPC3, when the Application Quality of Service rate-limiter is configured to specific traffic, packet loss occurs on unrelated traffic until reboot. PR1394085
- The **set security flow log dropped-illegal-packet** and **set security flow log dropped-icmp-packet** CLI commands are unhidden. PR1394720

Interfaces and Chassis

- On SRX4600 device, the virtual IP address of the VRRP might not respond to host-inbound traffic. PR1371516
- The following message appears for each port whose settings are changed or refreshed:

  Apr 3 12:00:00 srx /kernel: check_configured_tpids: <interface> : default tpid (0x8100) not configured. pic allows maximum of 0 tpids
  Apr 3 12:00:00 srx /kernel: check_configured_tpids: <interface> number of configured tpids exceeds the limit(0)

  PR1373668

Interfaces and Routing

- On SRX1500 devices, the ae0 and ae1 interfaces display the MAC address as 00:00:00:00:00:00 and 00:00:00:00:00:01. PR1352908

Intrusion Detection and Prevention (IDP)

- IDP signature update fails on the secondary node. PR1358489
- The IDP might not be deployed due to IDP configuration is not able to commit. PR1374079
- The unified policies configured with IDP, might not inspect arbitrary sessions, and marking them as **Not Interested** within the **show security idp counters flow** command. PR1385094

J-Web

- When the J-Web fails to get the resource information, the Routing Engine CPU usage shows 100% in resource utilization in the J-Web dashboard. PR1351416
- The J-Web setup wizard does not propagate the DHCP attributes from ISP to LAN. PR1370700
- The chassis cluster image does not displayed on the J-Web dashboard. PR1382219

Layer 2 Features

- The dcpfe/fxpc process might crash when you try to allocate large memory 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 Address Translation (NAT)

- Source NAT sessions might fail to create when the **port-overloading** or the **port-overloading-factor** is configured. PR1370279

Network Management and Monitoring

- With user firewall enabled and RG0 failover is being performed, eventd process core files are generated. PR1366120

Platform and Infrastructure

- On SRX1500 devices, when the power supply fails, the trap sent might contain incorrect information. PR1315937
- Frequency logs are displayed on the SRX5400, SRX5600, and SRX5800 devices when the IOC card has the same identifier as the SPC PIC card. PR1357913
- The SCP configuration backup fails even though the **/var/etc/ssh_known_hosts** has a proper fingerprint. PR1359424
- On SRX4600 devices, the show chassis fan show chassis environment command does not display any output. PR1363645
- On SRX1500 device, continues alarm on fan is observed. PR1367334
- Packet capture feature does not work after removing the sampling configuration. PR1370779
- On SRX Series devices in chassis cluster, the cold synchronization process might slow down when there are many packet forwarding engines (PFEs) installed on the device. PR1376172
- Junos upgrade might fail with the validate option after the **/cf/var/sw** directory is erroneously deleted. PR1384319

Routing Policy and Firewall Filters

- The TCP protocol ports 5800 and 5900 are added to junos-defaults to support VNC application. PR1333206
- The **show security policies detail** command output is modified to improve readability, particularly for unified policies. PR1338307
- On SRX Series devices, the nsd process might crash on the Packet Forwarding Engine with large-scale security policy configuration. PR1354576

- DDynamic application autocomplete support is not functional within the CLI for the **show security match-policies** command. PR1363908
- On SRX4100, SRX4200, SRX4600, SRX4800, when dynamic application is configured in security policy core files are observed on the PFE. As a workaround, do not configure dynamic application in security policy. PR1368762
- 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

- If family iso is enabled through the GREoIPSec (GRE over IPSec) tunnel, the vFPC stops working. PR1364624

Services Applications

- When modifying the ICAP configuration and the traffic passing through, the core file might generated. PR1389600
- Clearing the TCP session might not clear the redirect objects. PR1390835

Unified Threat Management (UTM)

- The default action of Web filtering does not work as expected. PR1365389
- When the server port is configured as 443, the displayed EWF server status is **UP**. PR1383695

VPNs

- During an RG0 failover in ISSU, when you use the rekeys, the iked core process file are generated. PR1340973
- Policy-based VPN is not working with the virtual router. PR1350123
- IPsec tunnel might not work when there are concurrent IKEv2 Phase 1 SA rekeys. PR1360968
- On SRX5600 AND SRX5800 devices, during VPN to AutoVPN configuration migration, traffic loss is observed. PR1362317
- On SRX Series devices in chassis cluster, when the VPN configuration size reaches the internal configuration processing chunk size, the VPN tunnels might not be configured successfully and the VPN tunnels might not come up after a reboot, upgrade, or restart ipsec-key-management. PR1376134
- Packet loss is observed in IPsec Z-mode scenario. PR1377266
- The kmd process might stop and cause VPN traffic outage after running the **show security ipsec next-hop-tunnels** command. PR1381868
- Adding or deleting site-to-site manual NHTB VPN tunnels to an existing st0 unit causes the existing manual NHTB VPN tunnels under the same st0 unit to flap. PR1382694

Resolved Issues: 18.2R1

Application Layer Gateways (ALGs)

- On SRX Series devices with SIP ALG enabled, the SIP ALG might drop SIP packets that have a **referred-by** or **referred-to header** field containing multiple header parameters. PR1328266
- SIP calls drop when the limit per SPU crosses 10,000 calls. PR1337549

Authentication and Access Control

- On SRX Series devices, the Packet Forwarding Engine might crash and a huge number of core files might be generated within a short time. PR1326677
- On SRX Series devices, incomplete Request Support Information (RSI) might be seen. PR1329967
- On SRX Series devices, the sessions might close because of the **idle Timeout junos-fwauth-adapter** logs. PR1330926
- Web authentication uses hard-coded three seconds timeout but in some scenarios the three seconds timeout is too short to complete a web authentication. Use the new CLI **set access firewall-authentication web-authentication timeout** command to configure web authentication timeout value. PR1339627

Chassis Clustering

- The device might stop forwarding traffic after RG1 failover from node0 to node1. PR1323024
- IP monitoring is not working as expected when one node is in secondary hold state and the primary node's priority is 0. PR1330821

Class of Service (CoS)

- Packets go out of order on SPC2 cards with IOC1 or FIOC cards. PR1339551

Flow-Based and Packet-Based Processing

- On SRX4600 devices, when you execute the **clear security flow session** command, time taken to clear the session depends on the total session number. For example, the clear session takes 9 minutes to clear 57M sessions. PR1308901
- Periodic PIM register loop is observed during switch failure. PR1316428
- The OSPF peers are unable to establish neighbors between the LT interfaces of the logical systems. PR1319859
- The IPv6 traffic does not work as expected on IOC3 with the services offloading (npcache) feature. PR1331401
- SSH to the loopback interface of SRX Series devices does not work properly when AppTrack is configured. PR1343736
- The flowd process might stop when SYN-proxy function is used. PR1343920
- SNMP MIB walk provides incorrect data counters for total current flow sessions. PR1344352

- The interface MAC limit configured under VLANS, which is in the range of the CLI guideline, does not take effect. PR1347245
- File download stops over a period of time when TCP proxy is activated through AV or Juniper Sky ATP PR1349351
- On SRX Series devices in a chassis cluster, if an IPv6 session is being closed and at the same time the related data-plane Redundancy Group (RG1+) failover occurs, this IPv6 session on the backup node might hang and cannot be cleared. PR1354448
- On SRX5000 line devices, when the IPsec performance acceleration feature is enabled, packets going in or out of a VPN tunnel are dropped. PR1357616

Intrusion Detection and Prevention (IDP)

- The control plane CPU usage is high when using IDP. PR1283379
- Loading IDP policy fails due to less available heap memory. PR1347821

J-Web

- J-Web does not display wizards on the dashboard. PR1330283
- When httpd process is not running, J-Web setup wizard does not work after you run the **request system zeroize** command, . PR1335561
- In J-Web you cannot delete dynamic VPN user configuration. PR1348705
- In J-Web menu security policies search button using Internet Explorer version 11 does not work. PR1352910
- The unsupported et and xe interface parameter for speed, link mode, and media type are removed from the **Configure>Interface>Ports** tab in J-Web. PR1355871

Layer 2 Ethernet Services

- The default gateway route might be lost after the failover of RG0 in a chassis cluster. PR1334016

Network Address Translation (NAT)

- Arena utilization on a FPC spikes and then resumes to a normal value. PR1336228

Platform and Infrastructure

- When you perform commits with apply-groups, VPN might flap. PR1242757
- The packet captured by datapath-debug on an IOC2 card might be truncated. PR1300351
- Inconsistent flow-control status on the reth interface is observed. PR1302293
- On SRX5000 line devices using DC PEM, the output of the **show chassis environment pem** and **show chassis power** commands shows incorrect DC input values. PR1323256
- On SRX5400, SRX5600, and SRX5800 devices, SPC2 XLP stops processing packets in the ingress direction after repeated RSI collections. PR1326584

- When Security-Intelligence is configured, IPFD CPU utilization might be higher than expected. PR1326644
- The log messages file contains the **node*.fpc*.pic* Status:1000 from if_np for ifl_copnfig op:2 for ifl :104** message. PR1333380
- Log message **No Port is enabled for FPC# on node0** is generated every 5 seconds. PR1335486
- On SRX4100 devices, interfaces are shown as half-duplex, but there is no impact on the traffic. PR1358066

Routing Policy and Firewall Filters

- Flowd process stops after configuring a huge number of custom applications. PR1347822
- On SRX Series devices, a large-scale commit, for example, a 70,000-lines security policy, might stop the nsd process on the Packet Forwarding Engine. PR1354576

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 The traces specific to received messages (BGP RECV traces) are logged correctly. PR1318830
- On SRX Series devices, dedicated BFD does not work. PR1347662

Unified Threat Management (UTM)

- The ISSU upgrade might fail due to the Packet Forwarding Engine generating a core file. PR1328665

VLAN Infrastructure

- On SRX Series devices in transparent mode, the flowd process might stop when matching the destination MAC. PR1355381

VPN

- IPsec traffic statistic counters return 32-bit values. PR1301688
- PKID syslog for key-pair deletion is required for conformance. PR1308364
- SNMP for jnxIpSecTunMonVpnName does not work. PR1330365
- The kmd process might generate core files when all VPNs are down. PR1336368
- All IPsec tunnels are in both active and inactive state. PR1348767
- S2S tunnels are not redistributed after IKE and IPsec are reactivated in a configuration. PR1354440
- The iked process might crash when IKE and IPsec SA rekey happens simultaneously PR1420762

SEE ALSO

New and Changed Features | 342

Documentation Updates

There are no errata or changes in Junos OS Release 18.2R3 documentation for the SRX Series.

SEE ALSO

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 might 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

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://pathfinder.juniper.net/feature-explorer/

SEE ALSO

- New and Changed Features | **342**
- Changes in Behavior and Syntax | **354**
- Known Behavior | **357**
- Known Issues | **360**
- Resolved Issues | **363**
- Documentation Updates | **380**
- Migration, Upgrade, and Downgrade Instructions | **380**

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

- **Feature Explorer**—Juniper Networks Feature Explorer helps you in exploring software feature information to find the right software release and product for your network. https://apps.juniper.net/feature-explorer/
- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved. prsearch.juniper.net.
- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms. apps.juniper.net/hct/home

  NOTE: 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.

- **Juniper Networks Compliance Advisor**—Review regulatory compliance information about Common Criteria, FIPS, Homologation, RoHS2, and USGv6 for Juniper Networks products. apps.juniper.net/compliance/.

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:

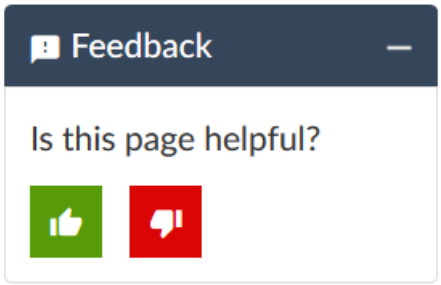

  - Click the thumbs-up icon if the information on the page was helpful to you.
  - Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- 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 postsales technical support, you can access our tools and resources online or open a case with JTAC.

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

Self-Help Online Tools and Resources

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

- Find CSC offerings: 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://casemanager.juniper.net/casemanager/

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) tool located at 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://casemanager.juniper.net/casemanager/.
- 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/cgi-bin/docbugreport/.

Revision History

28 January 2022—Revision 17, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 October 2021—Revision 16, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 August 2021—Revision 15, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 June 2021—Revision 14, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

05 February 2021—Revision 13, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

19 January 2021—Revision 12, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 11, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 November 2020—Revision 10, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 May 2020—Revision 9, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

26 March 2020—Revision 8, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

30 January 2020—Revision 7, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 December 2019—Revision 7, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 October 2019—Revision 6, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 October 2019—Revision 5, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

8 August 2019—Revision 4, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 July 2019—Revision 3, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 July 2019—Revision 2, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 June 2019—Revision 1, Junos OS Release 18.2R3– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

9 May 2019—Revision 13, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

19 April 2019—Revision 12, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

4 April 2019—Revision 11, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 March 2019—Revision 10, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 March 2019—Revision 9, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 February 2019—Revision 8, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 February 2019—Revision 7, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 January 2019—Revision 6, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

24 January 2019—Revision 5, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 January 2019—Revision 4, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 December 2018—Revision 3, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 December 2018—Revision 2, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 December 2018—Revision 1, Junos OS Release 18.2R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 October 2018—Revision 13, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 October 2018—Revision 12, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

20 September 2018—Revision 11, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

18 September 2018—Revision 10, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 September 2018—Revision 9, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 August 2018—Revision 8, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 August 2018—Revision 7, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 August 2018—Revision 6, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 August 2018—Revision 5, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

2 August 2018—Revision 4, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

19 July 2018—Revision 3, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

6 July 2018—Revision 2, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

29 June 2018—Revision 1, Junos OS Release 18.2R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.