Release Notes: Junos® OS Release 18.4R2 for the ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion

15 July 2021

Contents

Introduction

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

These release notes accompany Junos OS Release 18.4R2 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.

New Features in 18.4R2

Feature	Release Note Section
Support for 100-Mbps and 1-Gbps speeds Tri-Rate Copper SFP (ACX Series)	“New and Changed Features” on page 13
Transparent clock over IPv6 support (ACX Series)	“New and Changed Features” on page 13
Zero Touch Provisioning (ACX Series)	“New and Changed Features” on page 13
Layer 2 and Layer 3 VXLAN gateways (EX4650 and QFX5120 switches)	“New and Changed Features” on page 32, “New and Changed Features” on page 222
EVPN control plane and VXLAN data plane support (EX4650 and QFX5120 switches)	“New and Changed Features” on page 32, “New and Changed Features” on page 222
EVPN pure type-5 route support (EX4650 and QFX5120 switches)	“New and Changed Features” on page 32, “New and Changed Features” on page 222
Support for optimizing the SNMP walk execution time for IPsec statistics (MX Series)	“New and Changed Features” on page 90
Additional encapsulations added to pseudowire subscriber logical interfaces (MX Series)	“New and Changed Features” on page 90

Feature	Release Note Section
Assisted replication in data centers with EVPN-VXLAN overlay networks (QFX Series switches)	"New and Changed Features" on page 222
OVSDB support with VMware NSX for vSphere (QFX5120 switches)	"New and Changed Features" on page 222
Selective multicast forwarding support in EVPN-VXLAN (QFX Series)	"New and Changed Features" on page 222
BPDU protection in EVPN-VXLAN (QFX Series)	"New and Changed Features" on page 222

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 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 features and enhancements in Junos OS Release 18.4R2 for ACX Series Universal Metro Routers.

New and Changed Features: 18.4R2

Interfaces and Chassis

- **Support for 100-Mbps and 1-Gbps speeds on Tri-Rate Copper SFP (ACX5448)**—In Junos OS Release 18.4R2, ACX5448 routers support 100-Mbps and 1-Gbps speeds on Tri-Rate Copper SFP optics (740-013111). Note that 100-Mbps speed is supported only from ports xe-0/0/24 through xe-0/0/47.

  NOTE: 10-Mbps speed is not supported on Tri-Rate Copper SFP transceivers due to hardware limitations.

  - To set the speed for the optics, issue the **set interfaces *interface-name* speed auto** command. See speed for more details.
  - To enable autonegotiation, issue the **set interfaces *interface-name* gigether-options auto-negotiation** command. See auto-negotiation.

Software Installation and Upgrade

- **Zero Touch Provisioning (ACX5448)**—Starting in Junos OS Release 18.4R2, Zero Touch Provisioning (ZTP) automates the provisioning of the device configuration and software image with minimal manual intervention on management interface **em0**.

  When you physically connect a router to the network and boot it with a factory configuration, the router upgrades the Junos OS software image automatically and autoinstalls configuration file from the network through the management interface.

Timing and Synchronization

- **Transparent clock over IPv6 support (ACX5448)**—Starting with Junos OS Release 18.4R2, ACX5448 routers support transparent clock functionality for PTP over IPv6.

  To configure the transparent clock functionality, you must include the **e2e-transparent** statement at the **[edit protocol ptp]** hierarchy level.

  Use the **show ptp global-information** command to check the status of the transparent clock functionality configured on the router.

  [See Understanding Transparent Clocks in Precision Time Protocol]

New and Changed Features: 18.4R1

Authentication, Authorizing, and Accounting (AAA)

- **Support for password change policy enhancement (ACX Series)**—Starting in Junos OS Release 18.4R1, the Junos OS password change policy for local user accounts is enhanced to comply with additional password policies. As part of the policy improvement, you can configure the following:

- **maximum-lifetime-value**—The maximum duration of a password. The password expires after the maximum is reached.
- **minimum-lifetime-value**—The minimum duration of a password. You cannot change the password until the minimum duration is reached.

[See password.]

Interfaces and Chassis

- **Multichassis link aggregation groups, configuration synchronization, and configuration consistency check (MC-LAG) (ACX5448 routers)**—Starting in Junos OS Release 18.4R1, MC-LAG enables a client device to form a logical LAG interface using two switches. MC-LAG provides redundancy and load balancing between the two switches, multihoming support, and a loop-free Layer 2 network without running spanning tree protocols (STP).

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

MPLS

- **Support for topology independent loop-free alternate (TI-LFA) for IS-IS, advertising MPLS labels (ISIS, OSPF), and configuring SRGB for SPRING (ISIS, OSPF) (ACX5448)**—Starting with Junos OS Release 18.4R1, ACX5448 router support topology independent (TI)-loop-free alternate (LFA), advertise MPLS labels (ISIS, OSPF), and segment routing global block (SRGB) for SPRING (ISIS, OSPF).

  Topology independent (TI)-loop-free alternate (LFA) with segment routing provides fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure. You can enable TI-LFA for IS-IS by configuring **use-post-convergence-lfa** statement at the **[edit protocols isis backup-spf-options]** hierarchy level.

  You can configure SRGB range label used by source packet routing in networking (SPRING). The labels from this SRGB range is used for SPRING in IS-IS domain. This way the labels advertised in the segment routing is more predictable and deterministic across the segment routing domain.

  - To configure the starting index value of the SRGB label block, use the **start-label *start-label-block-value*** statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.
  - To configure the index range of the SRGB label block, use the **index-range *value*** configuration statement at the **[edit protocols isis source-packet-routing srgb]** hierarchy level.

  ACX5448 router supports IS-IS and OSPF segment routing enabled through MPLS. IS-IS and OSPF creates an adjacency segment per IS-IS and OSPF neighbor, for a given interface, adjacency, and area. A separate MPLS label is allocated for each adjacency segment created.

  To configure OSPF segment routing, use the following configuration statements at the **[edit protocols ospf]** hierarchy level:

  - **source-packet-routing**—Enable the source packet routing feature.
  - **node-segment**—Enable the node segment.

To configure IS-IS segment routing, use the following configuration statements at the **[edit protocols isis]** hierarchy level:

- **source-packet-routing**—Enable the source packet routing feature.
- **node-segment**—Enable source packet routing at all levels.

[See Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS, Understanding Source Packet Routing in Networking (SPRING), and source-packet-routing (Protocols IS-IS and OSPF).]

Platform and Infrastructure

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

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

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

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

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

Routing Protocols

- **Metro Ethernet services over segment routing infrastructure (ACX5448 routers)**—Starting with Junos OS Release 18.4R1, Metro Ethernet services are supported over a segment routing infrastructure.

  The following features are supported or can be configured:

  - IPv4 OSPF segment routing enabled through MPLS.
  - IS-IS segment routing enabled through MPLS.
  - Segment routing global block (SRGB) range label, which is used by Source Packet Routing in Networking (SPRING).
  - Anycast segment identifiers (SIDs) and prefix SIDs in SPRING are supported.
  - Topology independent loop-free alternate (TI-LFA) with segment routing, which provides fast reroute (FRR) backup paths corresponding to the post-convergence path for a given failure.

[See Understanding Adjacency Segments, Anycast Segments, and Configurable SRGB in SPRING for IS-IS Protocol, Understanding Topology-Independent Loop-Free Alternate with Segment Routing for IS-IS, Understanding Source Packet Routing in Networking (SPRING)].

Timing and Synchronization

- **Support for PTP boundary clocks for phase and time synchronization (ACX5448)**—Starting with Junos OS Release 18.4R1, ACX5448 routers support PTP boundary clocks for phase and time synchronization using IEEE-1588 Precision Timing Protocol (PTP). This feature also supports:
  - PTP over IPv4 (IEEE-1588v2)
  - PTP ordinary and boundary clocks
  - One-step clock mode operation for the PTP master clock
  - 10-MHz and 1-PPS output for measurement purpose

  All PTP packets use the best-effort queue instead of the network control queue.

  The ACX5448 router does not support the following features:

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

  [See IEEE 1588v2 PTP Boundary Clock Overview.]

VPNs

- **Support to control traceroute over Layer 3 VPN (ACX Series)**—Starting in Junos OS Release 18.4R1, in a Layer 3 VPN topology with **vrf-table-label** configured and multiple customer edge (CE) routers configured in the same VPN routing and forwarding (VRF) routing instance, when traceroute is performed to a remote provider edge (PE) router for a CE-facing network, the ICMP time exceeded packet determines the correct IP address as the source address.

  To control the traceroute over Layer 3 VPN topology with vrf-table-label configured and multiple CE routers configured in the same VRF, you can configure **allow-l3vpn-traceroute-src-select** at the **[edit system]** hierarchy level that determines the correct IP source address by reviewing the destination routing instance and destination IP address.

  [See allow-l3vpn-traceroute-src-select.]

SEE ALSO

Changes in Behavior and Syntax | 18

Known Behavior | 20

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 18.4R2 for the ACX Series routers.

Changes in Behavior and Syntax: 18.4R2

IN THIS SECTION

Interfaces and Chassis

- **Support for creating layer 2 logical interface independently (ACX Series)**—In Junos OS Releases 18.4R1, 18.4R2 and later, ACX Series routers support creating layer 2 logical interface independent of layer 2 routing instance type. That is, you can configure and commit the layer 2 logical interfaces separately and add the interface to bridge-domain or Ethernet VPN (EVPN) routing instance separately. Note that the layer 2 logical interfaces works fine only when the interface is added to bridge domain or EVPN routing instance.

  In the earlier Junos OS releases, when an layer 2 logical interface configuration (units with encapsulation vlan-bridge configuration) is used, then the logical interface must be added as part of a bridge-domain or EVPN routing instance for the commit to succeed.

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when change in number of supported history records is detected (ACX Series)**—In Junos OS Release 18.4R2, when Ethernet Connectivity Fault Management (ECFM) starts, it detects the number of history records supported by the existing Performance Monitoring history database and if there is any change from the number of history records supported (that is, 12) in 18.4R2, then the existing Performance Monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

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.]

Changes in Behavior and Syntax: 18.4R1

IN THIS SECTION

Network Management and Monitoring

- **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 18.4R1, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, the server must not return an RPC reply that encloses both an **<rpc-error>** element and an **<ok/>** element. If the operation is successful, but the server reply would enclose one or more **<rpc-error>** elements of severity warning in addition to the **<ok/>** element, then 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 encloses both an **<rpc-error>** element of severity warning and an **<ok/>** element.

SEE ALSO

New and Changed Features | 13

Known Behavior | 20

Known Issues | 22

Resolved Issues | 25

Documentation Updates | 29

Migration, Upgrade, and Downgrade Instructions | 30

Product Compatibility | 31

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 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

- When Layer 3 packets are classified, DiffServ code points are not preserved but are getting lost at the egress interface because of a chipset limitation.PR1322142
- ARP learning rate is very low. PR1343221
- All PTP packets goes to the best-effort queue instead of Network Control queue. This is because of the limitation on Qumran where DSCP values are not preserved" PR1361315
- PTX10001-20C Junos Telemetry Interface or Telemetry infrastructure does not support the interface-filtering capability. Therefore, after you enable a particular sensor for telemetry, it is turned-on for all the interfaces. PR1371996
- For et-interfaces, only PRE_FEC_SD defect is raised and no OTN alarm is raised. PR1371997
- The CLI **static-cak** command encryption does not work between two ACX-OX transpoder nodes. PR1389802
- For ACX6360 TIC beacon port-range needs to be updated to 0-7 instead of 0-15. PR1399335
- If user configures an invalid speed configuration on TIC ports (PIC slot 1) on ACX6360-OR/OX, the TIC interfaces are not created. PR1403546
- Junos OS do not perform vlan-id check at the egress and vlan-id check is only performed at ingress. PR1403730
- Single physical interface and logical interface for both ICCP and ICL is not a supported model on ACX5448 platform. Only in this model, static ARP configuration to peer IRB IP is required. PR1410971
- ACX5448 Forwarding-class defaults setup in schema file for ACX5448. Therefore, it is expected to see a firewall process **LIBCOS_COS_** errors are seen while upgrading the devices with latest software image/build when it is attempted to read pvidb db. PR1422284
- On Ethernet bridge, L2 filters might not work as expected when trying to match VLAN-based fields for untagged packets. PR1423214
- Some TPIDs (like TPID 9200) are not supported with VLAN-CCC and some of the other features . As these TPIDs are not supported rewrite is not possible on those TPIDs. PR1433500

Routing Protocols

- When multiple adjacencies are coming-up or flapping, some routes might not have remote-lfa backup next hops. They will appear only after next SPF trigger either manually or through network event. PR1389392

SEE ALSO

Known Issues

IN THIS SECTION

There are no known issues in hardware and software in Junos OS Release 18.4R2 for the ACX Series Router.

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

General Routing

- Forwarding when using non-existing SSM map source address in IGMPv3 instead of pruning. This is a day 1 design issue which needs to be redesigned. The impact is more, But definitely this needs some soaking time in DCB before it gets ported in previous versions. PR1126699
- When ACX2100 and ACX2200 routers are used as ingress PE routers for l2circuit connections, and the PE-CE interface (UNI) is an aggregated Ethernet interface, then upon MPLS path switchover, the traffic might silently get dropped or discarded. PR1194551
- The maximum number of logical interfaces on a ACX5000 Series has been increased from 1000 to 4000. PR1229492

- On ACX5448 routers when 1-Gigabit SFP is plugged in the router, autonegotiation is enabled by default. There is no functional impact. Only the CLI **show interfaces <intf-name> extensive** will show the autonegotiation field as disabled. PR1343679
- On ACX5000 platforms with Junos OS Release 16.2 and later, if the ECC errors occur, the FPC/fxpc process might use high CPU. This issue can be hit after the upgrade in some cases. PR1360452
- PTPIP packets are queuing to best effort queue instead of network control queue by default. PR1361315
- ARP request is getting dropped and not forwarded to the NNI interface queue when the CoS configuration has temporal buffer size. PR1363153
- On ACX Series routers running in Junos OS Release 15.1R1 to 15.1R8, when configuring **mac-table-size** under bridge-domain, an incorrect commit error appears not allowing the commit to pass. PR1364811
- Dedicated minimum buffers are reserved for some queues according to the Junos OS working model. These buffers are always available to those queues irrespective of the traffic pattern throughout the system. When the **clearing stat** statement is used, these values are visible. This cosmetic or minor issue has no functional impact. PR1367978
- SD threshold can be set above SF threshold. PR1376869
- Because of a race condition, in which the **class-of-service** configuration request for an interface is received before the e1-interface is created, a circuit with specified class-of-service parameters is created. Because of this, the interface creation fails, resulting in traffic not flowing on the e1-interface and then (if e1-interfaces are further disabled or enabled) a core file is generated. PR1378747
- Host bound traffic might be affected and lt interface might go down in ACX Series routers. PR1382166
- Enhancement is needed for FRR BER threshold SNMP support. PR1383303
- On ACX6360 and PTX10001-20C the Tx power cannot be configured using + sign. PR1383980
- Customer should avoid using the **loss-priority high** command in the firewall filters (MF classifiers). PR1388731
- On Junos OS Release 17.3 and later, the ACX5000 Packet Forwarding Engine syslog frequently shows the following errors messages: **acx_cos_tcp_bind_queues:736 parent acx_cos_tcp_ifd for ifd:ae0 doesn't exist for ifl:549**. In Junos OS Release 17.3R3-S1 the error logs appear only from time to time, and this can be related with an interface flap. In Junos OS Release 18.1R3, the logs appear constantly, without any interface flap. This message is related to HCOS checking (even without HCOS configured). 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
- The ccc logs are not compressed after rotation. PR1398511
- A jnxIfOtnOperState trap notification is sent for all ot-interfaces. This is a day-1 issue. PR1406758
- On ACX1000, ACX2000, ACX4000, ACX5048, and ACX5096 platforms, after a new child logical interface with VLAN and filter is added on an aggregated Ethernet physical interface or changing the VLAN ID of a child logical interface with filter, traffic over the aggregated 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

- The CFP2-DCO maximum achievable Tx power varies from module to module and with modulation format. The settable range in software is capped at 10 dBm. This limit is based on the class 1 laser safety limit, and not every module is able to achieve 10 dBm of Tx power. Tx power is set to the module maximum achievable power, if the user configured power is not reachable. This ensures that the module Tx power can be used up to its maximum power. PR1408194
- With ACX5448 platform devices, the ztp process will proceed with image upgrade even in situations when there is a mismatch in platform name of the software image stored on ftp/ztp servers and actual platform where the ztp process is being run. PR1418313
- On ACX5000 platform, the fxpc process high CPU usage might be seen under rare condition if parity errors are detected in devices. It has no direct service/traffic impact. However, since CPU utilization is high during this issue, there are some side-effects. For example, it might impact time-sensitive features like BFD. PR1419761
- On ACX5448 platforms, the optics might not work properly due to the fact that an internal defect causes it to not output power. As a result, the interface might not come up. PR1424814
- The Tx laser is enabled by default in CPLD. Therefore, the link shows up on the peer as soon as the Packet Forwarding Engine starts. PR1430910
- If auto-RP is used to signal RP, ACX5448 platform is periodically timeouting auto-RP mapping. That is because the proper port and queue towards CPU are not specified which causes RP announce message not reaching Routing Engine. Hence, RP mapping timeout. PR1432889
- Protocols get forwarded when using non-existing SSM map source address in IGMPv3 instead of pruning. PR1435648
- Commit fails while configuring a firewall policer action as "forwarding-class". PR1446556

Interfaces and Chassis

- 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 might prevent deletion of the IP avoiding the deletion of the family inet of the unnumbered interface. PR1412534

Layer 2 Features

- On Junos OS ACX5000, on the interfaces where lldp is 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 and kernel tries to remove the implicit filters, which return ENOENT, since entries were already disabled during the first commit. The following messages are harmless to the system. PR1400606

SEE ALSO

New and Changed Features | 13

Changes in Behavior and Syntax | 18

Known Behavior | 20

Resolved Issues | 25

Documentation Updates | 29

Migration, Upgrade, and Downgrade Instructions | 30

Product Compatibility | 31

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in Junos OS 18.4R1 for ACX 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:18.4R2

Class of Service (CoS)

- Error message **STUCK_BUFF : port_sp not empty for port 35 sp 1 pkts:1** is seen. 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
- On an ACX Series ring topology, after link between ACX Series and MX Series flap, VPLS RI on PE (MX) have no MAC of CE over l2 circuit. PR1360967
- On ACX5000, **fpc0 (acx_rt_ip_uc_lpm_install:LPM route add failed) Reason : Invalid parameter** error is seen after configuring **lpm-profile**. PR1365034
- On ACX5448, **LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** prints while commit on configuration prompt. PR1376665
- 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 flap. PR1381487
- On ACX 5448, 100-Gigabit link FEC enabled by default on 100-Gigabit LR4. PR1389518
- On ACX Series platforms the **forwarding-option dhcp-relay forward-only** configuration statement stops working and the DHCP packets are dropped. PR1392261
- On ACX Series, MTU is not properly applied and output of **ping mpls l2circuit sweep** is giving lower values than expected. PR1393947
- On ACX5048, the rpm rfc2544-benchmarking test fails to start. PR1395730
- Error message **ACX_PFE_ERROR: dnx_cfm_bd_endpoint_create: Failed to destroy the remote endpoint, Endpoint id 0x2001001, Entry not found** is logged. PR1397878
- CFM adjacency is not going down with distinct intervals. PR1397883
- Error message **ACX_ASIC_PROGRAMMING_ERROR: dnx_cfm_bd_endpoint_create: Failed to create the local endpoint Invalid parameter** been logged on peer node. PR1397951
- **Output packet error Count** incrementing on 100GE, 40GE ports on RIO. PR1398270
- High jsd or na-grpcd CPU usage might be seen even JET or JTI is not used. PR1398398
- Dynamic tunnels are not supported on ACX Series routers. PR1398729

- FPC might crash after offline/online MIC-3D-16CHE1-T1-CE-H. PR1402563
- VLAN tagged traffic arriving on VPLS interface might get dropped. PR1402626
- The ot/et interface is not created when invalid speed is configured. PR1403546
- On ACX 5448, TrTCM policer configuration parameters are as per RFC4115. PR1405798
- **show services inline stateful-firewall flow** or **show services inline stateful-firewall flow extensive**ccommand might cause the memory leak. PR1408982
- ACX Series routers drops DNS responses which contain an underscore. PR1410062
- The aggregated Ethernet interface Twamp history statistics verification on 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/PCCreate message if there is metric type other than type 2. PR1412659
- Number of inet-arp policers implemented on ACX 5000 is increased from 16 to 64. PR1413807
- The swap memory does not get initialized on boot on ACX5048. PR1415898
- Commit error while configuring firewall with term having log/syslog and accept actions. PR1417377
- On ACX5448 routers, BFD timer values are not as per the configured 900ms with multiplier 3, its showing 6.000 with multiplier 3 instead for most of the sessions. PR1418680
- CoS table error can sometimes cause traffic outages and SNMP timeouts if the optic is plugged out and inserted back. PR1418696
- Copying images from WAN interface to Routing Engine of ACX5448 takes long time. PR1422544
- The FPC or fxpc crash might be observed on an ACX Series platforms. PR1427362

Interfaces and Chassis

- Upgrading to Junos OS Release 17.4R1 results in generating cfmd process core file. PR1425804

MPLS

- MPLS ingress LSP's for LDP link protection are not coming up after disabling or enabling MPLS. PR1432138

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:18.4R1

General Routing

- Incorrect packet statistics are reported in the ifHCInUcastPkts OID. PR1306656
- ACX Series routers support from dual-tagged through untagged packets Layer 3 traffic. PR1307666
- Port xe-0/3/0 did not come up. PR1328207
- ACX Series routers are incorrectly allowing to configure higher values in **burst-size-limit** than what is supported by the hardware. PR1361482
- ACX Series routers autonegotiation shows incorrect values for link-partner when using SFP-LH or SFP-SX transceivers in combo-ports or SFP ports. PR1362490
- FEC PM error counters are accumulated instead of resetting after bin rollover. PR1363270
- VPLS with **vlan-id-list** is not working properly in some releases when the link between a PE device and a CE device is an aggregated Ethernet interface with a single member link and child physical interface flap. PR1365894
- The **commit** or **commit check** operation might fail because of the error **cannot have lsp-cleanup-timer without lsp-provisioning**. PR1368992
- The fxpc might crash after an interface is changed on ACX5000 routers. PR1378155
- Timestamp is incorrect for BER statistics after clearing. PR1386253
- The **request chassis beacon** CLI command is not working for pic-slot 1 (that is, CFP2 ports). PR1386711
- Certain builds of Junos OS do not allow you to upgrade or commit configuration changes when the SI service interface is used. PR1393729
- ACX Series routers does not support **physical-interface-filter** semantic in egress direction for any filters. It supports **interface-specific** command only. PR1395362
- High jsd or na-grpcd CPU usage might be seen when JET or JTI is not used. PR1398398

Platform and Infrastructure

- On Junos OS, the next-hop index allocation fails and private index space get exhausted through incoming ARP requests to management interface (CVE-2018-0063). PR1360039

SEE ALSO

New and Changed Features | 13

Changes in Behavior and Syntax | 18

Known Behavior | 20

Known Issues | 22

Documentation Updates | 29

Migration, Upgrade, and Downgrade Instructions | 30

Product Compatibility | 31

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 for the ACX Series documentation.

SEE ALSO

New and Changed Features | 13

Changes in Behavior and Syntax | 18

Known Behavior | 20

Known Issues | 22

Resolved Issues | 25

Migration, Upgrade, and Downgrade Instructions | 30

Product Compatibility | 31

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 | 13
- Changes in Behavior and Syntax | 18
- Known Behavior | 20

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for EX Series Switches

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 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 Junos OS Release 18.4 for the EX Series.

NOTE: The following EX Series switches are supported in Release 18.4R2: EX2300-C, EX3400, EX4300, EX4600-40F, EX4650, EX9200, EX9204, EX9208, EX9214, EX9251, and EX9253.

Release 18.4R2-S3 New and Changed Features

EVPNs

- **Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface (EX4600 switches)**—You can configure and commit the following on a physical interface of an EX4600 switch in an EVPN-VXLAN environment:
  - Layer 2 bridging (**family ethernet-switching**) on any logical interface unit number (unit 0 and any nonzero unit number).
  - VXLAN on any logical interface unit number (unit 0 and any nonzero unit number).
  - Layer 2 bridging (**family ethernet-switching** and **encapsulation vlan-bridge**) on different logical interfaces (unit 0 and any nonzero unit number).
  - Layer 3 IPv4 routing (**family inet**) and VXLAN on different logical interfaces (unit 0 and any nonzero unit number).

  For these configurations to be successfully committed and to work properly, you must specify the **encapsulation flexible-ethernet-services** configuration statement at the physical interface level—for example, **set interfaces xe-0/0/5 encapsulation flexible-ethernet-services**.

  [See Understanding Flexible Ethernet Services Support With EVPN-VXLAN.]

Release 18.4R2 New and Changed Features

EVPNs

- **Layer 2 and Layer 3 VXLAN gateways (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 18.4R2, you can deploy EX4650 and QFX5120 switches as follows:
  - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
  - (QFX5120 switches only) As a Layer 2 or Layer 3 VXLAN gateway in an Open vSwitch Database (OVSDB) overlay network

  VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See Understanding VXLANs.]

- **EVPN control plane and VXLAN data plane support (EX4650 and QFX5120 switches)—**Starting with Junos OS Release 18.4R2, EX4650 and and QFX5120 switches support EVPN-VXLAN. By using a Layer 3 IP-based underlay network coupled with an EVPN-VXLAN overlay network, you can place endpoints anywhere in the network and remain connected to the same logical Layer 2 network.

  EVPN-VXLAN is commonly deployed over the following physical underlay architectures:

  - A two-layer IP fabric that includes spine devices (Layer 3 VXLAN gateways) and leaf devices (Layer 2 VXLAN gateways). You can deploy EX4650 or QFX5120 switches as spine or leaf devices in this fabric.
  - An edge-routed briding overlay, which is a one-layer IP fabric that includes leaf devices that function as both Layer 2 and Layer 3 VXLAN gateways. You can deploy EX4650 or QFX5120 switches as leaf nodes in this fabric.

  [See Understanding EVPN with VXLAN Data Encapsulation.]

- **EVPN pure type-5 route support (EX4650 and QFX5120 switches)—**Starting with Junos OS Release 18.4R2, you can configure pure type-5 routing in an EVPN-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 ip-prefix-routes.]

Software Defined Networking

- **Layer 2 and Layer 3 VXLAN gateways (EX4650 and QFX5120 switches)—**Starting with Junos OS Release 18.4R2, you can deploy EX4650 and QFX5120 switches as follows:

  - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
  - (QFX5120 switches only) As a Layer 2 or Layer 3 VXLAN gateway in an Open vSwitch Database (OVSDB) overlay network

  VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

[See Understanding VXLANs.]

Release 18.4R1 New and Changed Features

Hardware

- **2-port QSFP+/1-port QSFP28 uplink module for EX4300-48MP and EX4300-48MP-S switches**—Starting with Junos OS Release 18.4R1, EX4300-48MP and EX4300-48MP-S switches support the 2-port QSFP+/1-port QSFP28 uplink module. The 2-port QSFP+/1-port QSFP28 uplink module can house two QSFP+ transceivers or one QSFP28 transceiver.

  [See EX4300 Switch Hardware Guide.]

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for password change policy enhancement (EX Series)**—Starting in Junos OS Release 18.4R1, the Junos OS password change policy for local user accounts is enhanced to comply with additional password policies. As part of the policy improvement, you can configure the following:

  - **maximum-lifetime-value**—The maximum duration of a password. The password expires after the maximum is reached.
  - **minimum-lifetime-value**—The minimum duration of a password. You cannot change the password until the minimum duration is reached.

  [See password.]

EVPNs

- **Support for graceful restart on EVPN-VXLAN (EX9200)**—Starting in Junos OS Release 18.4R1, Junos OS supports graceful restart on EVPN-VXLAN on EX9200 and QFX Series switches and MX Series Routers. Graceful restart allows the device to recover from a routing process restart or Routing Engine switchover without nonstop active routing (NSR) enabled.

  [See NSR and Unified ISSU Support for EVPN Overview.]

- **Support for VMTO for ingress traffic (EX9200)**—Starting in Junos OS Release 18.4R1, you can configure a leaf or spine device that is configured as a Layer 3 gateway to support virtual machine traffic optimization (VMTO) for ingress traffic. VMTO eliminates the unnecessary ingress routing to default gateways when a virtual machine is moved from one data center to another.

  To enable VMTO, configure **remote-ip-host** routes at the **[edit routing-instances routing-*instance-name* protocols evpn]** hierarchy level. You can also filter out the unwanted routes by configuring an import policy under the **remote-ip-host routes** option.

  [See Configuring EVPN Routing Instances.]

- **Support for multihomed proxy advertisement (EX9200)**—Starting in Junos OS Release 18.4R1, Junos OS now provides enhanced support to proxy advertise the MAC address and IP route entry from all leaf devices that are multihomed to a customer edge (CE) device. Using proxy advertisement prevents traffic

loss when one of the connections to the leaf device fails. To support the multihomed proxy advertisement, all multihomed provider edge (PE) devices should have the same multihomed proxy advertisement bit value. The multihomed proxy advertisement feature is enabled by default, and Junos OS uses the default multihomed proxy advertisement bit value of 0x20.

[See EVPN Multihoming Overview.]

- **MLD snooping support for EVPN-MPLS (EX9200)**—Starting with Junos OS Release 18.4R1, you can configure Multicast Listener Discovery (MLD) protocol snooping on EX9200 switches in an EVPN over an MPLS network. Enabling MLD snooping helps to constrain IPv6 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 to multiple PE devices.

  MLD snooping support in this environment includes:

  - Either MLDv1 and MLDv2 with any-source multicast (*,G) or MLDv2 with source-specific multicast (S,G) (configurable)
  - MLD state synchronization among multihoming PE devices using BGP EVPN Type 7 (Join Sync Route) and Type 8 (Leave Sync Route) network layer reachability information (NLRI)
  - Inclusive multicast forwarding from the ingress PE device into the EVPN core to reach all other PE devices
  - Forwarding across bridge domains (VLANs) using IRB interfaces and PIM operating in passive and distributed designated router (PIM-DDR) modes

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

Forwarding and Sampling

- **Support for activating or deactivating static routes on the basis of RPM test results (EX Series)** —Starting in Junos OS 18.4R1, you can use RPM probes to detect link status, and change the preferred-route state on the basis of the probe results. Tracked routes can be IPv4 or IPv6, and support a single IPv4 or IPv6 next hop. For example, you can send RPM probes to an IP address to determine whether the link is up, and if it is so, take the action of installing a static route in the route table. RPM-tracked routes are installed with preference 1 and are thus preferred over any existing static routes for the same prefix.

  [See Configuring RPM Probes , rpm-tracking, and show route rpm-tracking.]

Interfaces and Chassis

- **Support for uplink module with two 40-Gigabit Ethernet ports and one 100-Gigabit Ethernet port (EX4300-48MP)**—Starting with Junos OS Release 18.4R1, the 2-port QSFP+/1-port QSFP28 uplink module on EX4300-48MP switches can be configured to operate either two 40-Gigabit Ethernet ports or one 100-Gigabit Ethernet port. By default, the uplink module operates only the two 40-Gbps ports. To enable 100-Gbps speed, issue the **set chassis fpc 0 pic 2 port 0 speed 100g** command. The uplink module then enables the 100-Gigabit Ethernet port and disables the adjacent 40-Gigabit Ethernet ports.

  NOTE:
  - You can install the 2-port QSFP+/1-port QSFP28 uplink module only in PIC slot 2 on the switch.
  - You can configure 100-Gbps speed only on port 0 of PIC 2 (which is the uplink module slot on the switch).

  You can also channelize 40-Gigabit Ethernet interfaces, to four independent 10-Gigabit Ethernet interfaces using breakout cables.

  [See Setting the Mode on 2-port QSFP+/1-port QSFP28 Uplink Module (CLI Procedure).]

Junos Telemetry Interface

- **Packet Forwarding Engine and Routing Engine sensor support for Junos Telemetry Interface (JTI) (EX4600 switches)**—Starting in Junos OS Release 18.4R1, JTI supports Packet Forwarding Engine and Routing Engine statistics for EX4600 switches:

  The following Routing Engine statistics are supported through JTI:

  - LACP state export
  - Chassis environmentals export
  - Network discovery chassis and components
  - LLDP export and LLDP model
  - BGP peer information (RPD)

- RSVP interface export
- RPD task memory utilization export
- LSP event export
- Network Discovery ARP table state
- Network Discovery NDP table state

The following Packet Forwarding Engine statistics are supported through JTI:

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

To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

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

Multicast

- **Multicast VLAN registration (MVR) (EX2300 and EX3400 switches and Virtual Chassis)**—Starting in Junos OS Release 18.4R1, EX2300 and EX3400 switches and Virtual Chassis support multicast VLAN registration (MVR). MVR efficiently distributes IPTV multicast streams across an Ethernet ring-based Layer 2 network, reducing the bandwidth required for this traffic by using a multicast VLAN (M-VLAN) over which multicast traffic is forwarded to interested listeners on other VLANs that are configured as MVR receiver VLANs. You can configure MVR at the **[edit protocols igmp-snooping vlan *vlan-name* data-forwarding] source** and **receiver** hierarchy levels, and use the **show igmp snooping data-forwarding** CLI command to view configured M-VLAN and MVR receiver VLAN associations. **(The feature described above is documented but not supported on EX2300 and EX3400 switches and Virtual Chassis in Junos OS Release 18.4R1.)**

  [See Understanding Multicast VLAN Registration.]

Port Security

- **Support for DHCP snooping and other access port security features on private VLANs (EX2300 and EX3400 switches and Virtual Chassis)**—Starting in Junos OS Release 18.4R1, you can enable Dynamic

Host Configuration Protocol (DHCP) snooping for security purposes on access ports that are in a private VLAN (P-VLAN). 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. This 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. The option is used to insert information about the network location of the remote host into DHCPv6 packets.
- DHCPv6 option 18—Circuit ID option for DHCPv6. The option is used to insert information about the client port into DHCPv6 packets.
- DHCPv6 option 16—Vendor ID option for DHCPv6. The option is 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.]

- **Untrusted mode on trunk interfaces for DHCP snooping (EX2300, EX3400, EX4300 and EX4600 switches)—**Starting in Junos OS Release 18.4R1, you can configure a trunk interface as untrusted for DHCP security. Trunk interfaces in untrusted mode support DHCP snooping and DHCPv6 snooping, dynamic ARP inspection (DAI), and IPv6 neighbor discovery (ND) inspection.

[See Understanding Trusted and Untrusted Ports.]

Virtual Chassis

- **Virtual Chassis support (EX2300-24MP, EX2300-48MP)**—Starting in Junos OS Release 18.4R1, multigigabit EX2300 switches can be interconnected into a Virtual Chassis with other EX2300 model switches as follows:
  - Any combination of up to four EX2300-24MP, EX2300-48MP, EX2300, and EX2300-C switches is supported.
  - You do not need to set mixed mode.
  - Any models of EX2300 switches can be in the master or backup Routing Engine roles.
  - Any 10-Gbps uplink ports installed with SFP+ transceivers can be configured as Virtual Chassis ports (VCPs) to interconnect member switches.
  - Use the same steps as for configuring any other EX2300, EX3400, or EX4300 Virtual Chassis.

  [See Understanding EX Series Virtual Chassis.]

VPNs

- **Support to control traceroute over Layer 3 VPN (EX Series)**—Starting in Junos OS Release 18.4R1, in a Layer 3 VPN topology with **vrf-table-label** configured and multiple customer edge (CE) routers configured in the same VPN routing and forwarding (VRF) routing instance, when you perform traceroute to a remote provider edge (PE) router for a CE-facing network, the ICMP time exceeded packet determines the correct IP address as the source address.

  To control the traceroute, configure **allow-l3vpn-traceroute-src-select** at **[edit system]** hierarchy level This configuration determines the correct IP source address by reviewing the destination routing instance and destination IP address.

  [See allow-l3vpn-traceroute-src-select.]

SEE ALSO

Changes in Behavior and Syntax | 41

Known Behavior | 43

Known Issues | 45

Resolved Issues | 52

Documentation Updates | 59

Migration, Upgrade, and Downgrade Instructions | 60

Product Compatibility | 61

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 18.4R2 for the EX Series.

Release 18.4R2-S3 Changes in Behavior and Syntax

Platform and Infrastructure

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

Release 18.4R2-S2 Changes in Behavior and Syntax

Platforms and Infrastructure

- 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.]

Release 18.4R2 Changes in Behavior and Syntax

Interfaces and Chassis

- **No support for performance monitoring on aggregated Ethernet Interfaces (EX4300)**—Y.1731 performance monitoring (PM) over aggregated Ethernet Interfaces is not supported on EX4300 switches.

  [See sla-iterator-profile.]

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.]

Security

- **Syslog or log action on firewall drops packets (EX4600 switches)**—Starting in Junos OS Release 18.4R2, if you configure are syslog and log references to the actual action terms configured in a firewall filter, control packets and ICMP packets sent to the Routing Engine might be dropped.

Release 18.4R1 Changes in Behavior and Syntax

Interfaces and Chassis

- **Enhanced AC PEM in high-line power configuration supplies 2400 W power (EX9204)**—Starting in Junos OS Release 18.4R1, on EX9204 switches, the enhanced AC PEM in high-line power configuration provides a power output of 2400 W. On Junos OS versions prior to 18.4R1, the PEM provided only 2050 W of power output.

  [See show chassis power.]

- **Support for creating layer 2 logical interface independently (EX Series)**—In Junos OS Releases 18.4R1, 18.4R2, and later, EX Series switches support creating layer 2 logical interface independent of layer 2 routing instance type. That is, you can configure and commit the layer 2 logical interfaces separately and add the interface to bridge-domain or Ethernet VPN (EVPN) routing instance separately. Note that the layer 2 logical interfaces works fine only when the interface is added to bridge domain or EVPN routing instance.

  In the earlier Junos OS releases, when an layer 2 logical interface configuration (units with encapsulation vlan-bridge configuration) is used, then the logical interface must be added as part of a bridge-domain or EVPN routing instance for the commit to succeed.

Network Management and Monitoring

- **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 18.4R1, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, the server must not return an RPC reply that encloses both an **<rpc-error>** element and an **<ok/>** element. If the operation is successful, but the server reply would enclose one or more **<rpc-error>** elements of severity warning in addition to the **<ok/>** element, then 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 encloses both an **<rpc-error>** element of severity warning and an **<ok/>** element.

Security

- **Firewall warning message (EX2300 switches)**—Starting in 18.4R1, a warning message is displayed whenever a firewall term includes log or syslog with the accept filter action.

SEE ALSO

New and Changed Features | 32

Known Behavior | 43

Known Issues | 45

Resolved Issues | 52

Documentation Updates | 59

Migration, Upgrade, and Downgrade Instructions | 60

Product Compatibility | 61

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 for the EX Series.

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

Class of Service (CoS)

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

EVPN

- When a VLAN uses an IRB interface as the routing interface, the **vlan-id** parameter must be set to "none" to ensure proper traffic routing. This issue is platform independent. PR1287557

General Routing

- When vlan is added as an action for changing the VLAN in both ingress and egress filters, the filter will not be installed PR1362609
- A few error messages related to the function **rt_mesh_group_add_check()** are seen during reboot. These errors are harmless. PR1365049
- Automatic channelisation is not supported for 40GBASE-BXSR, QSFP+40GE-LX4, QSFP-100G-PSM4, and 100GBASE-BXSR optics. PR1366103
- On the EX4300-MP switch, the et-0/2/* (100-Gigabit Ethernet) interface multicast queue in strict-priority mode gets the priority treatment only across other multicast queues. PR1377692

Routing Protocols

- On EX4650 switches, 254 neighbors and 200,000 routes can be scaled for IS-IS v4. Beyond 200,000 routes with 254 neighbors, adjacency flaps and traffic drop will be seen. However, with 40 neighbors, scaling of 351,000 routes is achieved. PR1368106

Virtual Chassis

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

SEE ALSO

New and Changed Features | 32

Changes in Behavior and Syntax | 41

Known Issues | 45

Resolved Issues | 52

Documentation Updates | 59

Migration, Upgrade, and Downgrade Instructions | 60

Product Compatibility | 61

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.4R2 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

General Routing

- On EX4300 switches, when **storm-control** or **storm-control-profiles** with **action-shutdown** is configured, if the storm-triggered traffic is control traffic such as LACP, the physical interface will be put into an STP blocking state rather than turned down, so valid control traffic might be trapped to the control plane and unrelated interfaces might be set down as an LACP timeout. PR1130099
- When you issue the **request system reboot** command, ZTP is triggered. During the mounting stage, **/var/db/scripts/import** does not get created, which later causes the configuration to be committed partially. This is seen in the warning **Warning: Commit failed, activating partial configuration. Warning: Edit the router configuration to fix these errors**. We have not identified the root cause of the problem. PR1289782
- On EX Series and QFX Series switches that support Virtual Chassis, the commit warning message **interface matching is supported only in a stand-alone Device** is seen although it is on a standalone device. You might see the warning message after a commit operation with **from interface** condition in the firewall filter on a single device. PR1296767
- The EX9253 does not support interface ranges for channelized interfaces, so you need to configure the interfaces individually. PR1350635
- The working uplink module SFP-T might go down in Junos OS Release 17.2R1 and later. PR1360602
- On QFX5000 and EX4650 platforms, if lcmd is restarted, a chassisd core file is generated with traffic drop for few seconds. PR1363652
- The 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

- Multicast router advertisement (RA) packets coming on a VLAN need to be flooded to interfaces of all FPCs belonging to the same VLAN. When packets traverse HighGig ports (which connect different FPCs), they need to pass through the hardware filter to transmit packets to other FPCs. However, the filter is not applicable for the HighGig ports, so multicast RA packets do not traverse other FPCs. PR1370329
- When CoS-based forwarding (CoS) is enabled, due to the indexed next-hop installation issue in the kernel, the rpd process might crash upon route flap and LSP flap. PR1374558
- An EX4300 configured with a firewall filter on lo0 and DHCP security on the VLAN simultaneously might drop legitimate DHCP renew requests from clients on the corresponding VLANs. This occurs due to implementation design and Broadcom chipset limitation. PR1376454
- In EX2300 and EX3400 switches, image upgrade might fail due to the **insufficient space** issue. PR1376488
- When the **show** command is taking a long time to display results, the STP might change states as BPDUs are no longer processed and cause too many 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 may be dropped. PR1395186
- 1-Gbps speed configuration support on EX9251. PR1400651
- PXE installation might fail in this release due to a failure in image upgrade after PXE initialization. PR1406743
- On an EX9200 device with MC-LAG 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
- On QFX5110 and QFX5120 and EX4650 platforms, uRPF check in strict mode does not work properly. PR1417546
- Issue with installing EFL license on EX4300-XXMP devices only. When you add a license, it fails to be added, as shown in this example: **{master:0} user@host> request system license add terminal Mar 01 12:03:05 [Type ^D at a new line to end input, enter blank line between each license key] EmergencyJUNOS285602007 aeaqia qmlbjd amrrha 2tcmbr gayaqb ycsbdm mjggim gbastv nzuxaz lsebew 45dfoj xgc3ah fbo6ct 7vv3hl ykp4zq 5g6xch szi7aq 3pek5e vh4myw jdi5wq dxyi3c rkgydi 3crzkr szq terminal:1 error: EmergencyJUNOS285602007: license not valid for this product add license failed (1 errors)** This issue affects only EFL licenses (AFL is not affected) and EX4300-MP devices. In order to fix it, upgrade to Junos OS Release 18.3R3 and later or Release 18.4R2 and later.PR1421033
- The factory-default configuration for EX4300, EX2300, EX3400, and EX4300-MP platforms now includes DHCP client configuration on IRB and VME to facilitate connectivity to phone-home server (redirect.juniper.net) from phone-home client running on the device. The factory-default configuration includes the following. 1. DHCP enabled on VME and IRB 2. Default VLAN with VLAN ID 1 and l3-interface as irb.0 PR1423015
- On EX2300, EX3400, EX4300, EX4600, and on QFX Series except Virtual Chassis platforms, if **igmp-snooping** is enabled, multicast traffic might be dropped silently. PR1423556
- Multiple EX platforms might be unable to commit baseline configuration after zeroize.

{master:0}[edit] root# commit check Mar 26 05:50:48 mustd: UI_FILE_OPERATION_FAILED: File /var/run/db/enable-process.data doesn't exist Mar 26 05:50:48 mgd[1938]: UI_FILE_OPERATION_FAILED: Failed to open /var/run/db/enable-process.data+ file error: Failed to open /var/run/db/enable-process.data+ file error: configuration check-out failed: daemon file propagation failed. PR1426341

- In case the port connecting the server from client is a trunk port and has multiple VLANs configured then for VLAN on which NDI is not configured ,the client remains in solicit state on starting dhcpv6 device. PR1428769
- 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
- NDI cannot be used in VLAN with IRB on EX92XX: neighborhood advertisements/ solicit packets destined to host are getting dropped with NDI Inspection (under DHCPv6 security) on a VLAN with IRB configuration on EX92XX in Junos OS Releases 18.4 onward. PR1439844
- BUM traffic Rate limiting done after removing Ethernet headers. L1 TX rate on ingress interface: 1G Tx rate with headers: 865Mbps Rx rate on the egress interface:800M L1 RX rate on egress interface: 925Mbps Storm control functionalities in MX-L card is achieved by poilcer and hence the below mentioned policer inaccuracy is applicable for storm control feature as well. The below is mentioned in the Hyperion FS clearly. Policer inaccuracy ? Since XM sprays packets to 4 different LUs, each LU will be processing packets of varying sizes.XM does not do strict round-robin, so even if all the incoming packets were to be of exact same sizes (which is not a practical scenario), each LU will still be loaded differently, hence there will be some periods where some LUs policing limit may reach sooner than the others (either due to processing more packets or due to processing larger packets). Hence, it is possible that, some LUs, who see the policing limit reached sooner may drop the packet or color them differently that might result into eventual drop while the other LUs could queue the packets for transmission; We could see this behavior within a single flow as well. Hence the policier functionality can be unpredictable at times. In an extreme case, a packet flow may be sent to a single LU and the policer result is 1/4th of what it is expected. Since the policer functionality, in general, may not work correctly, we will see the impact on all the policing features e.g. input-policer, three-color-policer (srTCM, trTCM), output-policer. PR1442842
- The issue is limited to a database related to a MAC move scenario. When **dhcp-security** is configured, if MAC move happens for multiple IPv4 and IPv6 clients happens, the jdhcpd might consume 100 percent CPU and later crash. PR1425206

Infrastructure

- The DCD (Data Carrier Detect) modem control signal is not implemented in UART driver for EX3400 and EX2300 platforms. Hence, ?log-out-on-disconnect? feature will not be functional on these platforms. PR1351906
- Junos OS might hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on a Linux and QEMU hypervisor. PR1359339

- When an SNMP poll is performed for the following OIDs, the backup Routing Engine returns the value 6 (6=down) for the fan and 1 (1=unknown) for the PSUs, even though the fan and PSUs are up. Fan: 1.3.6.1.4.1.2636.3.1.13.1.6 PSU: 1.3.6.1.4.1.2636.3.1.13.1.6.2 For a permanent fix, upgrade the chassis to Junos OS Release 15.1R8 or later. PR1360962
- On an EX3400 or EX2300, during ZTP with configuration and image upgrade with FTP as file transfer, image upgrade is successful but sometimes a VM core file is observed. PR1377721
- In a PVLAN with 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 **inter-switch-link** configuration statement is used when a private VLAN (PVLAN) spans multiple switches. PR1388186
- On EX2300/EX2300-C/EX2300-MP platforms, if Junos software is with FreeBSD kernel version 11 with the build date on or after 2019-02-12, the switch may stop forwarding traffic or responding to console. A reboot is required to restore the service. PR1442376

Junos Fusion Enterprise

- Junos Fusion is not able to add new satellite devices when MC-LAG is configured on an EX Series. PR1374982
- On EX4300 when a 10-Gigabit Ethernet fiber port is using 1-gigabit SFP optics, autonegotiation is enabled by default. BCM recommendation is to disable the autonegotiation for PHY84756 ports to bring up the satellite device. PR1420343
- In Junos Fusion Enterprise environment with EX2300-48P or EX2300-48T acting as satellite devices, the loop-detect feature does not work for ports 0-23, because the loop-detect filter is not properly applied. PR1426757

Layer 2 Features

- The **eswd[1200]: ESWD_MAC_SMAC_BRIDGE_MAC_IDENTICAL: Bridge Address Add: XX:XX:db:2b:26:81 SMAC is equal to bridge mac hence don't learn** error message 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

Layer 2 Ethernet Services

- On QFX5000 series or EX4300, EX4600, EX2300, and EX3400 platforms with Spine-Leaf scenario, when some (two or more than two) underlay interfaces with ECMP are brought down on Leaf devices, the Multi-Hop BFD overlay sessions between spines and leafs might flap. If BFD flaps, the protocols depending on BFD (typically, IBGP Protocols) would also flap, which leads to traffic impact. PR1416941

Layer 3 Features

- From the code analysis, the CPU rate limiting and corresponding queue points to 100 pps in Junos OS Release 12.3 for ARP traffic. But in case of Junos OS Release 11.4, the rate limiter value was 3 Kpps. PR1165757

Multicast

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

Network Management and Monitoring

- In a rare case where trace files are not properly closed by Junos OS, traceoption logs might stop writing to a log file. PR1380764

Platform and Infrastructure

- On EX4300 or EX4300 Virtual Chassis, if VLAN Spanning Tree Protocol (VSTP) is configured, when some operations about VSTP (for example, deactivating/activating VSPT interface, deactivating/activating VSPT VLAN, and so on) are done, the pfex process might crash. PR1178539
- On an EX4300 PoE-supported platform with LLDP or LLDP-MED enabled, if the maximum power is set to 20 W (Class-4), the maximum power drawn by all ports might be 12.5W, and it might take about an hour for all the ports to come up. PR1253517

- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur. Even with enhanced convergence configured, there is no guarantee that sub-second convergence will be achieved. PR1371493
- ICMPv6 packets are hitting the dynamic ingress filter with the higher priority, thus never reaching an MF or static classifier. PR1388324
- Adding the second IRB configuration to an aggregated Ethernet and then removing it causes the first IRB to stop working. PR1423106

Routing Protocols

- On a dual Routing Engine system with GRES and graceful restart enabled, if Bidirectional Forwarding Detection (BFD) with the **hold-down-interval** option is enabled on an external BGP peer, this BGP peer might stay in idle state after a Routing Engine switchover. PR1324475
- On EX4300/EX4600/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 should not be dropped and reach the routing engine. PR1379718
- In a multicast routing scenario using PIM, if you are configuring a static route with qualified-next-hop for a multicast source, process the rpd might crash. This is because the qualified next hop points to the GF_DLI (Gateway Family Data Links) address that the PIM is unable to process, resulting in the crash. PR1408443
- The 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. This message may or may not impact any subscribers coming up. The earlier issue where few subscribers were seen offline along with this message is fixed by PR1417574 but the message is still seen. PR1439514

Subscriber Access Management

- The authd reuses addresses too quickly before jdhcpd completely cleans up the old subscriber which flooding error log. Log examples: **jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815**. PR1402653

SEE ALSO

New and Changed Features | 32

Changes in Behavior and Syntax | 41

Known Behavior | 43

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in the Junos OS Release 18.4R2 for EX Series switches.

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

Resolved Issues: 18.4R2

EVPN

- The device might proxy the ARP probe packets in an EVPN environment. PR1427109
- Configuring ESI on a single-homed 25G port might not work. PR1438227

General Routing

- On QFX5120 and EX4650, the convergence delay between PE1 and P router link is more than the expected delay value. PR1364244
- 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 can increase internal kernel memory usage. PR1369638
- RIPv2 update packets might not be sent with IGMP snooping enabled. PR1375332
- EX-4300 Virtual Chassis : Commit error is observed for the first time while loading the Mini-PDT base configurations. PR1383469
- On QFX5120 and EX4650, occasionally two of the channelized 25Gbps Ethernet ports using 4x25-Gigabit breakout cable do not come up after Junos OS reboot. PR1384898

- EX3400-Virtual Chassis: The **Error tvp_status_led_set** and **Error:tvp_optics_diag_eeprom_read** logs are seen. PR1389407
- The **Input rate pps** is not increased on EX2300-MP uplink ports if the packet is a pure Layer 2 packet such as non-etherII or non-EtherSnap. PR1389908
- Interface flapping on an EX3400 Virtual Chassis causes 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 might be dropped if IGMP and PTP TC are configured together. PR1395186
- On EX2300 the MAC table is not populated after the **interface-mode** value is changed. PR1396422
- The fxpc core file might be seen if scaled number of filter-based forwarding (FBF) filters are configured. PR1398256
- High jsd or na-grpcd CPU usage might be seen when JET or JTI is not used. PR1398398
- EX3400 might not learn 30,000 MAC addresses while sending MAC learning traffic. PR1399575
- MAC limit with persistent MAC does not after reboot. PR1400507
- The authd process might crash when you issue **show network-access requests pending** command during the restarting of authd. PR1401249
- The TCP connection between ppmd and ppman might be dropped due to a kernel issue. PR1401507
- The **adt7470_set_pwm** message is continuously getting displayed after upgrade to Junos OS Release 18.1R3.3. PR1401709
- The STP does not work when the aggregated Ethernet interfaces number is AE1000 or above in QFX5000 and AE480 or greater in other QFX Series 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
- EX4300-48MP: Packets are dropped after the traffic filter and routing instance are configured. PR1407424
- 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
- The l2cpd might crash if the **vstp traceoptions** and **vstp-vlan-all** commands are configured. PR1407469
- EX3400 PSU status continues to be **check** even though the 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
- On EX Series and QFX Series switches, PEM Alarm for the backup FPC remains on Master FPC although backup FPC was detached from the Virtual Chassis. PR1412429
- On EX4300-48MP, the chassis Status LED shows yellow instead of amber. PR1413194
- chassisd output power budget received continually every 5 seconds without any alarm after upgrade to Junos OS 18.1R3 PR1414267

- VXLAN Encapsulation next hop (VENH) doesnt get installed during BGP flap or restart routing. PR1415450
- On EX3400, the **show chassis environment** repeats **OK** and **Failed** at short intervals. PR1417839
- The EX3400 Virtual Chassis status might be unstable during the bootup of the Virtual Chassis or after the Virtual Chassis port flaps. PR1418490
- Virtual Chassis might become unstable and FXPC core files when there are multiple configured filter entries. PR1422132
- On EX3400, autonegotiation status shows incomplete on ge-0/2/0 using SFP-SX. PR1423469
- 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
- Incorrect model information while polling through SNMP from Virtual Chassis. PR1431135

Infrastructure

- IfSpeed and IfHighSpeed erroneously reported as zero on EX2300. PR1326902
- The Packet Forwarding Engine is flooded with messages: **pkt rx on ifd NULL unit 0** PR1381151
- The dot1x could not work when dot1x is configured with isolated VLAN on one interface. PR1404664

Interfaces and Chassis

- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. PR1402606
- The IFLs in EVPN routing instances might flap after committing configurations. PR1425339

Junos Fusion Enterprise

- PoE over LLDP negotiation is not supported on a Junos Fusion Enterprise setup. PR1366106
- **error: peer_daemon: bad daemon: scpd** error is seen on EX9251 running Junos OS Releases 18.1R1 and 18.1R2. PR1369646
- Juniper Fusion Enterprise: Cannot log in to SD cluster although it is recognized by AD properly. PR1395570
- The l2ald might crash if you issue the **clear ethernet-switching table persistent-learning** command. PR1409403
- Extended ports in Junos Fusion Enterprise do not adjust the MTU value when VoIP is enabled. PR1411179
- Traffic might get discarded silently in a Junos Fusion Enterprise scenario with dual aggregation devices. PR1417139

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 you issue the **clear ethernet-switching table persistent-learning** command. PR1381739

Layer 2 Ethernet Services

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

Network Management and Monitoring

- Overtemperature trap is not sent out even though there is a Temperature Hot alarm. PR1412161

Platform and Infrastructure

- Ping does not go through device after WTR timer expires in ERPS scenario. PR1132770
- EX4300 upgrade fails during validation of the SLAX script during upgrade. PR1376750
- ECMP route installation failure with log messages such as unilist install failure might be observed on the EX4300 switch. PR1376804
- Unicast DHCP request get misforwarded to backup RTG link on EX4300 Virtual Chassis. PR1388211
- Continuous log messages get displayed on EX4300 after upgrading to a Junos OS Release 17.4 or later release. PR1391942
- EX4300 OAM LFM might not work on an extended VLAN bridge interface with native VLAN configured. PR1399864
- Traffic drop is seen on EX4300 when the 10-Gigabit Ethernet fiber port is using 1-gigabit Ethernet SFP optics with autonegotiation enabled. PR1405168
- The policer might not work when it is applied through the dynamic filter. PR1410973
- EX4300 QinQ - untagged UNI Traffic egress as single-tagged on NNI Interface. PR1413700
- EX4300 does not send **fragmentation needed** message when MTU is exceeded with DF bit set. PR1419893
- The traffic to the NLB server might not be forwarded if the NLB cluster works in multicast mode. PR1411549
- The pfex process might crash and core files generated when a SFP transceiver is reinserted. PR1421257
- Traffic might be lost when one of the logical interfaces on the LAG is deactivated or deleted. PR1422920
- The authd process crashes when the Accounting RADIUS server is not reachable. PR1424030

- EX9200-12QS switch sends tagged packets through the access interface and through the trunk interface with a native VLAN ID. PR1424174
- Interface flapping scenario might lead to ECMP next hop install failure on EX4300s. PR1426760
- VIP might not forward the traffic if VRRP is configured on an aggregated Ethernet interface. PR1428124
- The ERPS failover does not work as expected on EX4300 device. PR1432397

Routing Protocols

- 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 a Virtual Chassis if IGMP snooping is configured. PR1413543
- The QFX Series and EX Series switch might not install all IRB MAC addresses when the device is initialized. PR1416025
- Sometimes, IGMP snooping might not work. As a workaround, restart the multicast-snooping process. PR1420921

Subscriber Access Management

- EX4300 **/var** file is showing full as the **var/log/dfcd_enc** file grows in size. PR1425000

Resolved Issues: 18.4R1

General Routing

- On the EX4300-32F, the MACsec session stays down on 1-Gigabit and 10-Gigabit Ethernet links after certain events, when events are performed with traffic running. PR1299484
- On EX2300 and EX3400 switches, the bridge ID is assigned to **02:00:00:00:00:10** irrespective of the base-MAC addresses. PR1315633
- Incorrect value of optical power is displayed. PR1326642
- On EX3400 and EX2300 switches, 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 Q-in-Q VLAN is added to or deleted from an interface on EX2300 or EX3400 switches. PR1334850
- Consideration of relaxing P-VLAN conflict rules during VLAN change for reauthentication and CoA scenarios. PR1346936
- The 40-Gigabit Ethernet interfaces might not forward traffic. PR1349675
- On EX2300, EX3400, and EX4300MP switches 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

- The traffic uses the original IRB MAC address if you are configuring a MAC address for an IRB interface. PR1359816
- On EX2300MP switches, the fan count is wrong in jnxFruName,jnxFilledDescr and jnxContainersCount.4. PR1361025
- The EX4300-MP MACsec AES-GCM-128-XPN and AES-GCM-256-XPN cipher suites are not supported for MGE ports. PR1362035
- FPM board status is missing in the SNMP MIB walk result. PR1364246
- The l2cpd process might crash when you configure MVRP with private VLAN and RSTP **interface-all**. PR1365937
- Virtual Chassis split followed by generation of fxpc core files might occur when VLAN members are scaled. PR1369678
- Unicast ARP packet loop might be observed in a DAI scenario. PR1370607
- NTP broadcast packets are not forwarded out on Layer 2 ports. PR1371035
- MAC refresh packet might not be sent out from the new primary link after an RTG failover. PR1372999
- BOOTP packets might be dropped if **BOOTP-support** is not enabled at the global level. PR1373807
- FPC might crash when the output interface flaps with analyzer or sampling configured. PR1374861
- The port access list group is not properly reallocating the TCAM slices. PR1375022
- The interface AE480 or above might be in STP discarding state on EX9200 switches. PR1378272
- On EX4300-48MP, the IP transit traffic hits the lo0 filter. PR1379328
- All interfaces belonging to a certain FPC might be lost after multiple GRES in Virtual Chassis. PR1379790
- The 802.1X configuration does not work with Microsoft NPS server. PR1381017
- On EX4300-48MP, as the session-option configuration under the access profile hierarchy is not applicable for EX Series and QFX Series, do not use that statement and options under it PR1385229
- On EX9200, a warning message **prefer-status-control-active is used with status-control standby** is seen whenever you commit a configuration. PR1386479
- On an EX2300 with Q-in-Q (**flexible-vlan-tagging**), you are unable to obtain the DHCP IP for the IRB interface after power-cycling the device. PR1387039
- The smid process might generate core files during sanity script execution on QFX5100 and EX4300. PR1391909

EVPN

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

High Availability (HA) and Resiliency

- The backup Routing Engine might go to database prompt after performing configurations such as remove and restore are performed. PR1269383

Infrastructure

- Core files might be generated upon attempt to commit a configuration. PR1376362

Junos Fusion Enterprise

- The **peer_daemon: bad daemon: scpd** error message is seen on EX9251 running Junos OS Releases 18.1R1 and 18.1R2. PR1369646

Layer 2 Features

- The firewall filter might not work correctly with the match condition of **dot1q-tag** on an EX Series switch. PR1369592
- RTG MAC refresh packets are sent out from non-RTG ports if the RTG interface belonging to the Virtual Chassis master flaps. PR1389695

Network Management and Monitoring

- On EX4600 switches, unsupported CLI configurations and show commands from the **cfm** hierarchy or sub-hierarchy are allowed. PR1359052
- While toggling multiple times between baseline and CFM configurations, all 30 CFM sessions are not up. PR1360907
- The event-policy generated traps are sent with UTC, even though the time zone is defined under the system hierarchy. PR1380777

Platform and Infrastructure

- Interface flapping is seen on an EX4300 switch. PR1361483
- Some interfaces cannot be added under the MSTP configuration. PR1363625
- On EX4300 and EX4600 switches, the l2ald process might crash in an 802.1x scenario. PR1363964
- The Packet Forwarding Engine might crash if frequent MAC moves are encountered. PR1367141
- The LLDP TLV with the wrong switch port capabilities might be sent. PR1372966
- Login lockout might never expire because the timestamps of **Lockout start** and **Lockout end** are same. PR1373803
- On EX4300-48MP, unsupported 1-gigabit optics in the 10-gigabit uplink module might cause interface traffic to be dropped. PR1374390

- Traffic might be silently discarded with indirect next hop and load balancing. PR1376057
- The IRB interface does not go down when the master 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
- The EX4300 device chooses a wrong bridge ID as the RSTP Bridge ID. PR1383356
- On EX4300-48MP mixed Virtual Chassis, the Power over Ethernet interface maximum power configuration on a member EX4300 gives an error if the power is configured to be more than 30 W. PR1383717
- Layer 3 IP route is destroyed after the Layer 2 next hop is changed. PR1389688

Routing Protocols

- On EX4300-48MP, stale VLAN entries might be seen after a script involving split or merge reboots is run continuously. PR1363739

SEE ALSO

New and Changed Features | **32**

Changes in Behavior and Syntax | **41**

Known Behavior | **43**

Known Issues | **45**

Documentation Updates | **59**

Migration, Upgrade, and Downgrade Instructions | **60**

Product Compatibility | **61**

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 documentation for the EX Series switches.

SEE ALSO

New and Changed Features | **32**

Changes in Behavior and Syntax | **41**

Known Behavior | **43**

Known Issues | **45**

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.

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

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for Junos Fusion Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 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

There are no new features in Junos OS Release 18.4R2 for Junos Fusion Enterprise.

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

SEE ALSO

- Changes in Behavior and Syntax | 63
- Known Behavior | 64
- Known Issues | 64
- Resolved Issues | 65
- Documentation Updates | 66
- Migration, Upgrade, and Downgrade Instructions | 67
- Product Compatibility | 72

Changes in Behavior and Syntax

There are no changes in behavior of Junos OS features and changes in the syntax of Junos OS statements and commands in Junos OS Release 18.4R2 for Junos Fusion Enterprise.

SEE ALSO

- New and Changed Features | 62
- Known Behavior | 64
- Known Issues | 64
- Resolved Issues | 65
- Documentation Updates | 66
- Migration, Upgrade, and Downgrade Instructions | 67
- Product Compatibility | 72

Known Behavior

There are no known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 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.

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.4R2 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

- Junos Fusion is not able to add new satellite devices when MC-LAG is configured on EX platform. PR1374982

- On EX4300 when 10G Fibre port is using 1G Ethernet SFP optics, Auto Negotiation is enabled by default. BCM recommendation is to disable the auto-negotiation for PHY84756 ports to bring up the satellite device. PR1420343
- 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

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in Junos OS Release 18.4R2.

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

Resolved issues: Release 18.4R2

Junos Fusion Enterprise

- PoE over LLDP negotiation is not supported on Junos Fusion Enterprise. PR1366106
- error: peer_daemon: bad daemon: scpd on EX9251 running 18.1R1 and 18.1R2 PR1369646
- Juniper Fusion Enterprise: Cannot login to satellite device cluster though it is recognized by the aggregation device. PR1395570
- The l2ald process might generate a core file if the **clear ethernet-switching table persistent-learning** command is issued. PR1409403
- Extended ports in JFE do not adjust MTU when VoIP is enabled. PR1411179
- The traffic might get dropped in a Junos Fusion Enterprise with dual aggregation devices. PR1417139

Resolved issues: Release 18.4R1

- In a Junos Fusion Enterprise, the scpd process does not run on the EX9251. PR1369646

SEE ALSO

New and Changed Features | 62

Changes in Behavior and Syntax | 63

Known Behavior | 64

Known Issues | 64

Documentation Updates | 66

Migration, Upgrade, and Downgrade Instructions | 67

Product Compatibility | 72

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 for Junos Fusion Enterprise documentation.

SEE ALSO

New and Changed Features | 62

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-x86-64-18.3B1.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-x86-64-18.3B1.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 switch running Junos OS can be converted only to SNOS 3.1 and later.
> - Either the switch must be set to factory-default configuration by using the **request system zeroize** command, or the following command must be included in the configuration: **set chassis auto-satellite-conversion**.

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

If you need to convert a satellite device to a standalone device, you must install a new Junos OS software package on the satellite device and remove 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 about EEOL releases and to review a list of EEOL releases, see https://www.juniper.net/support/eol/junos.html

Downgrading from Junos OS

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

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

To downgrade a Junos Fusion Enterprise from Junos OS Release 18.3R1, follow the procedure for upgrading, but replace the 18.3 **junos-install** package with one that corresponds to the appropriate release.

SEE ALSO

- New and Changed Features | 62
- Changes in Behavior and Syntax | 63
- Known Behavior | 64
- Known Issues | 64
- Resolved Issues | 65
- Documentation Updates | 66
- Product Compatibility | 72

Product Compatibility

IN THIS SECTION

- Hardware and Software Compatibility | 73
- Hardware Compatibility Tool | 73

Hardware and Software Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

- New and Changed Features | 62
- Changes in Behavior and Syntax | 63
- Known Behavior | 64
- Known Issues | 64
- Resolved Issues | 65
- Documentation Updates | 66
- Migration, Upgrade, and Downgrade Instructions | 67

Junos OS Release Notes for Junos Fusion Provider Edge

IN THIS SECTION

- New and Changed Features | 74
- Changes in Behavior and Syntax | 76
- Known Behavior | 77
- Known Issues | 77

These release notes accompany Junos OS Release 18.4R2 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 18.4R2 New and Changed Features

There are no new features or enhancements to existing features for Junos Fusion Provider Edge in Junos OS Release 18.4R2.

Release 18.4R1 New and Changed Features

Class of Service (CoS)

- **CoS support for Broadband Edge Subscriber Management functionality on Junos Fusion Provider Edge**—Starting with Junos OS 18.4R1, standard CoS feature support is provided for broadband edge subscriber management functionality on Junos Fusion Provider Edge, including classifiers and rewrite rules for subscriber interfaces and up to four levels of hierarchical scheduling, depending on hardware used.

  [See Understanding CoS on an Aggregation Device in Junos Fusion Provider Edge.]

Junos Fusion

- **Support for broadband edge subscriber management (Junos Fusion Provider Edge)**—Starting in Junos OS Release 18.4R1, Junos Fusion Provider Edge supports broadband edge subscriber management where the aggregation device functions as the broadband network gateway (BNG). The aggregation device is used as a single point of management to provision and manage the broadband services on the extended ports on the satellite devices. The extended ports function as access ports on the BNG and are connected to customer premise equipment.

  [See Broadband on Junos Fusion and Junos OS Broadband Subscriber Management and Services Library.]

- **Connectivity fault management (Junos Fusion Provider Edge)**—Starting in Junos OS Release 18.4R1, Junos Fusion Provider Edge supports distributed and inline connectivity fault management (CFM) on the extended ports on the satellite devices. The aggregation device initiates and processes the continuity check messages (CCMs) that are sent and received on the extended ports on the satellite devices. This feature supports CCMs for multiple up MEPs, Ethernet loopback and linktrace for a MEP, and delay measurement and synthetic loss measurement for performance monitoring between two MEPs.

  [See Connectivity Fault Management in Junos Fusion.]

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 main release and the maintenance releases for Junos Fusion Provider Edge.

Release 18.4R2 Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 18.4R2.

Release 18.4R1 Changes in Behavior and Syntax

There are no changes in default behavior and syntax for Junos Fusion Provider Edge in Junos OS Release 18.4R1.

SEE ALSO

- New and Changed Features | 74
- Known Behavior | 77
- Known Issues | 77
- Resolved Issues | 78
- Documentation Updates | 79
- Migration, Upgrade, and Downgrade Instructions | 80
- Product Compatibility | 89

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 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 Provider Edge

- Mix of virtual gateway configurations on irb associated with vlan for VRRP and VXLAN together are not supported in QFX10000, even if the configurations exists vxlan configurations takes precedence. PR1413878

SEE ALSO

New and Changed Features | 74
Changes in Behavior and Syntax | 76
Known Issues | 77
Resolved Issues | 78
Documentation Updates | 79
Migration, Upgrade, and Downgrade Instructions | 80
Product Compatibility | 89

Known Issues

There are no known issues in the Junos OS Release 18.4R2 for Junos Fusion Provider Edge.

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

SEE ALSO

Resolved Issues

IN THIS SECTION

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

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

Resolved Issues: 18.4R2

Junos Fusion Provider Edge

- BUM traffic might get dropped on peer Fusion Aggregation Device when the link between Satellite Device and local Aggregate Device goes down. PR1384440
- QFX5110 : auto-negotiation is not disabled in hardware after setting no-auto-negotiation option in CLI. PR1411852

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) goes down. PR1397992

Resolved Issues: 18.4R1

Junos Fusion Provider Edge

- In a Junos Fusion, the aggregation device LAG interface might flap during satellite device upgrade or downgrade. PR1321575
- 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
- The spmd core process might generate a core file after the **request support information** command is executed on the aggregation device. PR1375732

Junos Fusion Satellite Software

- The shutdown of the cascade port might lead to the invalidation of the MPC. PR1360876
- QFX satellite device might restart in Junos OS Fusion solutions when copper SFP is used. PR1369062

SEE ALSO

New and Changed Features | **74**

Changes in Behavior and Syntax | **76**

Known Behavior | **77**

Known Issues | **77**

Documentation Updates | **79**

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

Product Compatibility | **89**

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 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.R1 is different from that for earlier Junos OS releases.

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

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

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

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

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

5. Select the **Software** tab.

6. Select the software package for the release.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

10. Install the new **jinstall** package on the aggregation device.

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

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

- For 64-bit software:

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

  ```
  user@host> request system software add validate reboot
  source/jinstall64-18.R1.SPIN-domestic-signed.tgz
  ```

- For 32-bit software:

  ```
  user@host> request system software add validate reboot
  source/jinstall-18.R1.SPIN-domestic-signed.tgz
  ```

All other customers, use the following commands.

- For 64-bit software:

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

  ```
  user@host> request system software add validate reboot
  source/jinstall64-18.R1.SPIN-export-signed.tgz
  ```

- For 32-bit software:

  ```
  user@host> request system software add validate reboot
  source/jinstall-18.R1.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.R **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos Fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite

device, you first need to upgrade the target satellite device to an interim Junos OS software version that can be converted to satellite software. For satellite device hardware and software requirements, see Understanding Junos Fusion Software and Hardware Requirements

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

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

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

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

2. Clear the device:

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

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

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

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

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

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

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

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

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

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

Converting a Satellite Device to a Standalone Device

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

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

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

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

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

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

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

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

5. Review and accept the End User License Agreement.

6. Download the software to a local host.

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

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

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

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

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

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

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

9. Commit the configuration.

   To commit the configuration to both Routing Engines:

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

Otherwise, commit the configuration to a single Routing Engine:

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

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

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

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

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

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

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

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

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

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

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

Upgrading an Aggregation Device

When you upgrade an aggregation device to Junos OS Release 18.R, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

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

Downgrading from Junos OS Release 18.

To downgrade from Release 18. to another supported release, follow the procedure for upgrading, but replace the 18. **jinstall** package with one that corresponds to the appropriate release.

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

For more information, see the Installation and Upgrade Guide.

SEE ALSO

- New and Changed Features | **74**
- Changes in Behavior and Syntax | **76**
- Known Behavior | **77**
- Known Issues | **77**
- Resolved Issues | **78**
- Documentation Updates | **79**
- Product Compatibility | **89**

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

- New and Changed Features | 74
- Changes in Behavior and Syntax | 76
- Known Behavior | 77
- Known Issues | 77
- Resolved Issues | 78
- Documentation Updates | 79
- Migration, Upgrade, and Downgrade Instructions | 80

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

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 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 MX Series.

Release 18.4R2 New and Changed Features

Network Management and Monitoring

- **Support for optimizing the SNMP walk execution time for IPsec statistics (MX Series)**—Starting in Junos OS Release 19.2R1, you can optimize the SNMP walk execution time for IPsec statistics. To achieve this optimization, increase the cache lifetime of the IPsec related information (for example statistics and SA information) so that a single SNMP walk request is served for N number of IPsec Security Associations (SAs) with N number of queries made to the service PIC. IPsec statistics are now fetched by the burst mode, thereby reducing the load on the Routing Engine daemon, kmd. For different scale needs, we may have to tweak the hidden SNMP knob parameters, for example, with Dead Peer detection (DPD) having more number of tunnels without traffic and simultaneous SNMP walks.

Subscriber Management and Services

- **Additional encapsulations added to pseudowire subscriber logical interfaces (MX Series with MPC and MIC)**—Currently, the supported encapsulation type on the pseudowire subscriber interfaces include:

  - **Transport logical interfaces**—Circuit cross-connect (CCC) encapsulation.
  - **Service logical interfaces**:
    - Ethernet VPLS encapsulation
    - VLAN bridge encapsulation
    - VLAN VPLS encapsulation

  Starting in Junos OS Release 18.4R2, in addition to the existing encapsulation types, the following support is provided:

  - **Transport logical interfaces**—Ethernet VPLS encapsulation, and provision for terminating the interface on the l2backhaul-vpn routing-instance.
  - **Service logical interfaces**—Circuit cross-connect (CCC) encapsulation, and provision for terminating the interface on locally switched Layer 2 circuits.

  [See Pseudowire Subscriber Logical Interfaces Overview.]

Release 18.4R1 New and Changed Features

Hardware

- **Smart SFP and smart SFP+ support (MX Series)**—Starting in Junos OS Release 18.4R1, the smart SFP transceivers and smart SFP+ transceiver in Table 1 on page 92 and Table 2 on page 92 are supported on the listed MX Series routers.

Table 1: SFP Transceiver Support on the MX Series

SFP Model	Supported MPCs, MICs, and Platforms
SFP-GE-TDM-T1 SFP-GE-TDM-DS3 SFP-GE-TDM-E1 SFP-GE-TDM-STM1 SFP-GE-TDM-STM4	Supported MPCs: • MX-MPC1E-3D (with MIC) • MX-MPC1E-3D-Q (with MIC) • MX-MPC2E-3D (with MIC) • MX-MPC2E-3D-Q (with MIC) • MX-MPC2E-3D-NG (with MIC) • MX-MPC3E-3D-NG (with MIC)
	Supported MICs: • MIC-3D-20GE-SFP • MIC-3D-20GE-SFP-E • MIC-MACSEC-20GE
	Supported platforms: • MX80 (with MIC) • MX104 (fixed interfaces as well as MIC) • MX240, MX480, and MX960 (with MPC+ MIC)

Table 2: SFP+ Transceiver Support on the MX Series

SFP+ Model	Supported MPCs, MICs, and Platforms
SFPP-XGE-TDM-STM16	Supported MPCs: • MX-MPC1E-3D (with MIC) • MX-MPC1E-3D-Q (with MIC) • MX-MPC2E-3D (with MIC) • MX-MPC2E-3D-Q (with MIC) • MX-MPC2E-3D-NG (with MIC) • MX-MPC3E-3D-NG (with MIC)
	Supported MICs: • MIC-MACSEC-20GE
	Supported platforms: • MX80 (with MIC) • MX104 (fixed interfaces as well as MIC) • MX240, MX480, and MX960 (with MPC+ MIC)

See the [Hardware Compatibility Tool].

- **Support for 40-Gbps ports to operate at 1-Gbps or 10-Gbps speed (MX10008)**—Starting in Junos OS Release 18.4R1, you can use the Mellanox pluggable adapter (QSFP+ to SFP+ adapter or QSA; model number: MAM1Q00A-QSA) to convert quad-lane based ports to a single-lane based SFP+ port. The QSA adapter has the QSFP+ form factor with a receptacle for the SFP+ module. Use the QSA adapter to convert a 40-gigabit port to a 1-Gbps or a 10-Gbps port. You can plug-in a 10-Gbps SFP+ transceiver into the QSA adapter, which is inserted into the QSFP or QSFP+ ports of the MX10K-LC2101 line cards of the MX10008 router.

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for password change policy enhancement (MX Series)**—Starting in Junos OS Release 18.4R1, the Junos OS password change policy for local user accounts is enhanced to comply with additional password policies. As part of the policy improvement, you can configure the following:
  - **maximum-lifetime-value**—The maximum duration of a password. The password expires after the maximum is reached.
  - **minimum-lifetime-value**—The minimum duration of a password. You cannot change the password until the minimum duration is reached.

  [See password.]

Class of Service (CoS)

- **Support for five-level hierarchical CoS with dynamic interface set over dynamic interface sets (MX Series)** — Starting in Junos OS Release 18.4R1, five-level hierarchical CoS with the ability to configure dynamic interface sets over dynamic interface sets is supported on NG-MPC2E, NG-MPC3E, MPC5, and MPC7 line cards.

  [See stacked-interface-set (Dynamic Profiles).]

- **Support for dynamic and static logical interfaces in the same dynamic interface set (MX Series)** — Starting in Junos OS Release 18.4R1, you can apply dynamic and static logical interfaces in the same dynamic interface set on all MPCs that support four-level and five-level hierarchical CoS.

  [See Understanding Hierarchical CoS for Subscriber Interfaces.]

EVPN

- **Support for VMTO for ingress traffic (MX Series)**—Starting in Junos OS Release 18.4R1, you can configure a leaf or spine device that is configured as a Layer 3 gateway to support virtual machine traffic optimization (VMTO) for ingress traffic. VMTO eliminates the unnecessary ingress routing to default gateways when a virtual machine is moved from one data center to another.

  To enable VMTO, configure **remote-ip-host** routes at the **[edit routing-instances *routing-instance-name* protocols evpn]** hierarchy level. You can also filter out the unwanted routes by configuring an import policy under the **remote-ip-host routes** option.

  [See Ingress Virtual Machine Traffic Optimization.]

- **Support for multihomed proxy advertisement (MX Series)**—Starting in Junos OS Release 18.4R1, Junos OS now provides enhanced support to proxy advertise the MAC address and IP route entry from all leaf devices that are multihomed to a CE device. This can prevent traffic loss when one of the connections to the leaf device fail. To support the multihomed proxy advertisement, all multihomed PE devices should have the same multihomed proxy advertisement bit value. The multihomed proxy advertisement feature is enabled by default, and Junos OS uses the default multihomed proxy advertisement bit value of 0x20.

  [See EVPN Multihoming Overview.]

- **Automatically generated and assigned Ethernet segment identifiers in EVPN-VXLAN and EVPN-MPLS Networks (MX240, MX480, QFX5100, and QFX5110)—**Starting in Junos OS Release 18.4R1, you can configure aggregated Ethernet interfaces and aggregated Ethernet logical interfaces on which LACP is enabled to automatically generate and assign Ethernet segment identifiers (ESIs) to themselves. We support this feature in the following environments:

  - On MX240 or MX480 routers that are multihomed in active-standby or active-active mode in an EVPN-MPLS network.
  - On QFX5100 or QFX5110 switches that are multihomed in active-active mode in an EVPN-VLAN network.

- **MLD snooping support for EVPN-MPLS (MX Series and vMX)**—Starting with Junos OS Release 18.4R1, you can configure Multicast Listener Discovery (MLD) protocol snooping on MX Series routers with MPCs and vMX routers in an EVPN over an MPLS network. Enabling MLD snooping helps to constrain IPv6 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 to multiple PE devices.

  MLD snooping support in this environment includes:

  - Either MLDv1 and MLDv2 with any-source multicast (*,G) or MLDv2 with source-specific multicast (S,G) (configurable)
  - MLD state synchronization among multihoming PE devices using BGP EVPN Type 7 (Join Sync Route) and Type 8 (Leave Sync Route) network layer reachability information (NLRI)

- Inclusive multicast forwarding from the ingress PE device into the EVPN core to reach all other PE devices
- Forwarding across bridge domains (VLANs) using IRB interfaces and PIM operating in passive and distributed designated router (PIM-DDR) modes

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

- **Support for graceful restart on EVPN-VXLAN (MX Series)**—Starting in Junos OS Release 18.4R1, Junos OS supports graceful restart on EVPN-VXLAN on EX9200 and QFX Series switches and MX Series Routers. Graceful restart allows the device to recover from a routing process restart or Routing Engine switchover without nonstop active routing (NSR) enabled.

  [See NSR and Unified ISSU Support for EVPN Overview.]

Forwarding and Sampling

- **Support for activating or deactivating static routes on the basis of RPM test results (MX Series)**—Starting in Junos OS 18.4R1, you can use RPM probes to detect link status, and change the preferred-route state on the basis of the probe results. Tracked routes can be IPv4 or IPv6, and support a single IPv4 or IPv6 next hop. For example, RPM probes can be sent to an IP address to determine if the link is up, and if so, take the action of installing a static route in the route table . RPM-tracked routes are installed with preference 1 and thus are preferred over any existing static routes for the same prefix.

  [See Configuring RPM Probes , rpm-tracking, and show route rpm-tracking.]

General Routing

- **Avoid jlock hogs by configuring jlock hold time (MX Series)**—Starting with Junos OS Release 18.4R1, users can configure a jlock hold time threshold value via sysctl. This helps avoid jlock hogs (tight loops) in **ifd_walk** by dropping the jlock after the threshold time is reached. The default hold time is 50ms.

  [See sysctl() Function]

High Availability (HA) and Resiliency

- **BFD Client for segment routing (MX Series)**—This feature is not supported on Junos OS Release 18.4R1. You can configure Junos OS to run Seamless Bidirectional Forwarding Detection (S-BFD) over non colored segment routing tunnels and use S-BFD as a fast mechanism to detect path failures. You can configure **bfd-liveness-detection** at the **[edit protocols source-packet-routing segment-list]** hierarchy level for enabling path-level S-BFD for a segment list.

  [See Understanding Bidirectional Forwarding Detection (BFD).]

Interfaces and Chassis

- **Support for enhanced Switch Control Board (MX240, MX480, and MX960)**—Starting in Release 18.4R1, Junos OS supports the Enhanced Switch Control Board SCBE3-MX (model number: SCBE3-MX-S) on the MX240, MX480, and MX960 routers. The SCBE3-MX-S supports a pluggable Routing Engine and provides a control plane and data plane interconnect to each line card slot. The SCBE3-MX provides a fabric bandwidth of up to 480Gbps, using four fabric planes (with MPC7 line cards).

The following Routing Engines are supported on SCBE3-MX: RE-S-1800x2, RE-S-1800x4, RE-S-X6-64G, and RE-S-X6-128G.

The SCBE3-MX interoperates with the following existing line cards: MS-MPC, MPC2-NG, MPC3, MPC3-NG, MPC4, MPC5, and MPC7.

SCBE3-MX supports fabric hardening. It supports configuration of per fpc **bandwidth-degradation** and per fpc **blackhole-action**.

The SCBE3-MX does not interoperate with any previous-generation SCBs (SCB, SCBE, and SCBE2). Also, the SCBE3-MX does not support smooth upgrade.

[See SCBE3-MX Description]

- **VRF-aware syslog client (MX Series)**—Starting in Junos OS Release 18.4R1, the system log (syslog) client is completely VRF aware. If a server is reachable through a virtual routing and forwarding (VRF) instance, the syslog client can send log messages to the server. To specify the routing instance through which the remote server is reachable, use the **routing-instance** statement (introduced at appropriate hierarchies).

  In previous releases, the syslog client could send log messages to a server reachable through a VRF instance only if the server could be looked up using the default (inet.0 or inet6.0) routing table. If you set the **management-instance** statement, the server was reachable through that VRF instance but the syslog client could not send syslog messages to the server.

  [See Management Interface in a Non-Default Instance and routing-instance (Syslog).]

- **Layer 2 and Layer 3 protocols, platforms, and service features supported on MX10008**— Starting in Junos OS Release 18.4R1, MX10008 routers support the following features:
  - SFLOW—Overview of sFlow Technology
  - Inline Active Flow Monitoring—Understanding Inline Active Flow Monitoring and bridge-template
  - Two-Way Active Management Protocol (TWAMP)—See Understanding Two-Way Active Measurement Protocol on Routers
  - MPLS—MPLS Overview
  - RSVP—RSVP Overview
  - MPC—MX Series MPC Overview
  - IPv4, IPv6, OSPF, and BGP—IPv6 Overview, Understanding IPv4 Addressing, OSPF Overview, and Understanding BGP.
  - Network Time Protocol (NTP)—NTP Overview
  - IGMP Snooping—IGMP Snooping Overview
  - BGP persistence for IPv4 and IPv6 and Segregation between interface specific code and DCD core code—Understanding the Long-Lived BGP Graceful Restart Capability and dcd
  - Connectivity Fault Management (CFM)—Ethernet OAM Connectivity Fault Management
  - Integrated Routing and Bridging (IRB)—Understanding Integrated Routing and Bridging

- gnMI—Enabling “ON CHANGE” Sensor Support Through Network Management Interface (gNMI)
- Rewrite of the first three bits of IPv6 DSCP value—inet6-precedence (CoS Rewrite Rules)
- NSR—Nonstop Active Routing Concepts
- TACACS+ Authentication and TACACS+ System Accounting— Configuring TACACS+ Authentication and Configuring TACACS+ System Accounting

Junos Telemetry Interface

- **Export of subscriber accounting and dynamic interface and interface-set queue statistics through Junos Telemetry Interface (JTI) (MX Series Routers)** —Starting in Junos OS Release 18.4R1, you can export statistics associated with dynamic subscriber interface stacking through remote procedure calls (gRPC). Accurate statistics (actual transit statistics) sensor for the subscriber interface includes IP (total) and IPv6 ingress and egress packets and bytes. Queue statistics for dynamic interface and interface sets include include counts of transmitted and dropped packets and bytes. The queue statistics sensors are maintained per contributing slot (as in the case with AE). Separate metadata sensors convey more contextual information about the dynamic interface and interface sets are available. The metadata sensors are also eligible for ON_CHANGE streaming.

  To enable subscriber and queue statistics for telemetry, include the **subscriber-statistics** and **queue-statistics** statements at the **[edit dynamic-profiles *profile-name* telemetry]** hierarchy level.

  [See dynamic-profiles and Guidelines for gRPC Sensors (Junos Telemetry Interface).]

- **Expanded ON_CHANGE support for Junos Telemetry Interface (JTI) (MX960, MX2010, MX2020, PTX5000, PTX1000, and PTX10000)**—Starting in Junos OS Release 18.4R1, OpenConfig support through remote procedure call (gRPC) and JTI is extended to support additional ON_CHANGE sensors.

  Periodical streaming of OpenConfig operational states and counters collects information at regular intervals. ON_CHANGE support streams operational states as events (only when there is a change), and is preferred over periodic streaming for time-sensitive missions.

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

  - **/components/component**
  - **/components/component/name/**
  - **/components/component/state/type**
  - **/components/component/state/id**
  - **/components/component/state/description**
  - **/components/component/state/serial-no**
  - **/components/component/state/part-no**

  ON_CHANGE notification will be supported on all the hardware components displayed in the Junos OS CLI operational mode command **show chassis hardware**.

To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. To enable ON_CHANGE support, configure the sample frequency in the subscription as zero.

[See Guidelines for gRPC Sensors (Junos Telemetry Interface) and show chassis hardware.]

- **Support for NTF agent (MX240, MX480, MX960, MX2010, MX2020, PTX1000, PTX5000, PTX10000, and VMX)**—Junos OS exposes telemetry data over gRPC and UDP as part of the Junos Telemetry Interface (JTI). One way to stream JTI data into your existing telemetry and analytics infrastructure requires managing an external entity to convert the data into a compatible format. Starting in Junos OS Release 18.4R1, the NTF agent feature provides an on-box solution that allows you to configure and customize to which endpoint (such as IPFIX and Kafka) the JTI data is delivered and in which format (such as AVRO, JSON, and MessagePack) the data is encoded.

  [See NTF Agent Overview.]

- **Abstracted fabric interface support on Junos Telemetry Interface (JTI) (MX480, MX960, MX2008, MX2010, MX2020, and MX-ELM)**—Starting in Junos OS Release 18.4R1, JTI sensor support is available for abstracted fabric interfaces. An abstracted fabric interface is a pseudointerface that represents a first class Ethernet interface behavior. This sensor is only supported for node virtualization configurations on MX routers with an abstract fabric Interface as the connecting link between guest network functions (GNFs). JTI sensors will report interface-specific load-balancing and fabric queue statistics. They also will report aggregated statistics across all abstracted fabric interfaces hosted on a source Packet Forwarding Engine of local guest network functions (GNFs) along with the fabric statistics for all traffic ingressing from and egressing to the fabric from that Packet Forwarding Engine.

  JTI sensor support is for both gRPC sensors and native (UDP) sensors. Use the following resource path to configure JTI sensors:

  - **/junos/system/linecard/node-slicing/af-fab-stats/**

  To provision the sensor to export data through gRPC, use the **telemetrySubcribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig for Junos OS module. Starting in Junos OS Release 18.3R1, OpenConfig and Network Agent packages are bundled into the Junos OS image by default. Both packages support the Junos Telemetry Interface (JTI).

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

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

- **Enhanced IS-IS sensor support for Junos Telemetry Interface (JTI) (MX960, MX2020, PTX5000, PTX1000, and PTX10000)**—Starting in Junos OS Release 18.4R1, JTI supports OpenConfig Version v0.3.3 (from v0.2.1) for resource paths related to IS-IS link-state database (LSDB) streaming. The difference between the two versions results in changes, additions, deletions, or non-support for leaf devices related to the following IS-IS type length value (TLV) parameters and IS-IS areas:

- TLV 135: extended-ipv4-reachability
- TLV 236: ipv6-reachability
- TLV 22: extended-is-reachability
- TLV 242: router-capabilities
- IS-IS interface attributes
- IS-IS adjacency attributes

To provision the sensor to export data through gRPC streaming, use the **telemetry Subscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig and Network Agent packages, both of which are bundled into the Junos image in a default package named **junos-openconfig**.

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

Layer 2 VPN

- **Group VPN on AMS interface (MX Series)—**Starting in Junos OS Release 18.4R1, Junos OS supports load-balancing Group VPN services on AMS interfaces. AMS interfaces are a bundle of interfaces that function as a single interface and can be configured to load-balance traffic among the group members. To configure load balancing of Group VPN services on AMS interfaces, include the **ipsec-group-vpn** in the **[edit services service-set *service-set-name*]** hierarchy level to configure the service set and the **load-balancing-option** statements in the **service-interface** hierarchy of the AMS interface to enable load balancing.

  For more information on configuring AMS interfaces, see Configuring Aggregated Multiservices Interfaces.

  [See Group VPN on AMS Interfaces.]

MPLS

- **Track IGP metric for install prefixes (MX Series)**—Starting in Junos OS Release 18.4R1, you can let the install prefixes follow the metric of their corresponding IGP prefix so that the various RSVP protocol routes installed for the LSP can now each have their indivdual metric value. The **install-prefix** IGP metric tracking feature can be configured for all LSPs at the **[edit protocols mpls]** level or on a per-LSP basis at the **[edit protocols mpls label-switched-path]** hierarchy level.

  [See Install Prefix IGP Overview.]

- **Support for IP-based filtering and port mirroring of MPLS traffic (MX Series with MPC and MIC)**—Starting in Junos OS Release 18.4R1, you can apply inbound and outbound filters for MPLS family based on MPLS-tagged IPv4 and IPv6 parameters using inner payload match conditions, and enable selective port mirroring of MPLS traffic unto a monitoring device.

  To enable IP-based filtering, additional match conditions, such as IPv4 and IPv6 source and destination addresses, protocol, source and destination ports, and IPv4 and IPv6 source and destination prefix list, are added under the MPLS filter term **from** parameter.

To enable port mirroring, additional actions, such as **port-mirror** and **port-mirror-instance**, are added for all the match conditions under the filter term **then** parameter.

[See Understanding IP-Based Filtering and Selective Port Mirroring of MPLS Traffic.]

- **Static egress LSP with IPv6 next-hop**—Starting in Junos OS Release 18.4R1, you can configure static LSP on the egress router with the IPv6 as a nexthop address to forward IPv6 traffic. Static LSP supports nexthop indirection and link protection.

  [See Configuring Static Label Switched Paths for MPLS.]

Network Management and Monitoring

- **New major alarms on MX Series routers with MPC1 and MPC2**—Starting in Junos OS Release 18.4R1, 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”; for example:

  ```
  user@host> show chassis alarms

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

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

  You can check the syslog messages to learn 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/cm/ or call 1-888-314-JTAC (within the United States) or 1-408-745-9500 (from outside the United States).

Operation, Administration, and Maintenance (OAM)

- **Support for inline link fault management (MX Series)**—Starting in Junos OS Release 18.4R1, Junos OS supports inline mode for OAM link fault management (LFM) on MX Series routers. Inline LFM delegates the transmission and receipt of LFM keepalive packets from the periodic packet management (**ppm**) process on the line card to the forwarding ASIC (that is, to the hardware). Inline LFM reduces the load on the ppm process and can support LFM in-service software upgrade (ISSU) for non-Juniper peers (for a keepalive interval of 1 second). You can enable inline LFM by including the **hardware-assisted-keepalives** configuration statement at the **[edit protocols oam ethernet link-fault-management]** hierarchy level. To disable inline LFM, delete the **hardware-assisted-keepalives** statement. The **show oam ethernet link-fault-management detail** command displays the keepalive packet statistics. Starting from Release 18.4R1, when inline LFM is enabled, the keepalive packet statistics are not updated. In earlier releases, the **show oam ethernet link-fault-management detail** command displayed the keepalive packet statistics.

  [See Enabling Inline Transmission of LInk Fault Management Keepalives for Maximum Scaling.]

Routing Policy and Firewall Filters

- **Support for next-filter as a firewall filter action (MX Series)—**Starting in Junos OS Release 18.4R1, firewall filters can be configured to execute a sequence of firewall *filter* actions. The new **next-filter** option allows you to deploy a filter list and run a series of filters, similar to what is already available with **next-term** actions, and provides filter scale optimization. Up to eight filters can be chained in this way. The feature is not supported on logical systems, or on loopback and pseudo-interfaces.

  You can use a filter list to implement a mix of multifield-classification and firewall filter rules. For example, the first filter in the list can be used to perform a generic filter classification, and the subsequent filters can then do the actual filtering.

  [See input-chain and output-chain.]

- **Filter-based GRE encapsulation (MX Series)**—Starting in Junos OS Release 18.4R1, you can use **tunnel-end-point** commands to enable line-rate, filter-based, GRE tunneling of IPv4 and IPv6 payloads across IPv4 networks.

  This GRE encapsulation is not supported for logical systems or for MPLS traffic, and the route lookup for GRE encapsulated traffic is supported on the default routing instance only.

  The following commands are introduced for this feature:

  set firewall tunnel-end-point *tunnel-name* gre

  set firewall tunnel-end-point *tunnel-name* ipv4

  set firewall tunnel-end-point *tunnel-name* ipv6

  [See tunnel-end-point and Filter-Based Tunneling Across IPv4 Networks.]

Routing Protocols

- **Support for BGP flowspec redirect to IP (MX Series)**—Starting in Junos OS Release 18.4R1, BGP flow specification as described in BGP Flow-Spec Internet draft draft-ietf-idr-flowspec-redirect-ip-02.txt, *Redirect to IP Action* is supported. Redirect to IP action uses extended BGP community to provide traffic filtering options for DDoS mitigation in service provider networks. Legacy flow specification, as specified in the Internet draft draft-ietf-idr-flowspec-redirect-ip-00.txt, *BGP Flow-Spec Extended Community for Traffic Redirect to IP Next Hop*, redirect to IP uses the BGP nexthop attribute to support interoperability of devices. Junos OS advertises redirect to IP flow specification action using the extended community by default. Redirect to IP action allows you to divert matching flow specification traffic to a globally reachable address. This feature is required to support service chaining in virtual service control gateway (vSCG).

  To configure a static IPv4 flow specification route, include the **redirect *ipv4-address*** statement at the **[edit routing-options flow route then]** hierarchy level in the configuration.

  To configure a static IPv6 specification route, include the **redirect *ipv6-address*** statement at the **[edit routing-options flow route then]** hierarchy level in the configuration.

  To configure legacy flow specification include **legacy-redirect-ip-action** at the **[edit group *bgp-group* neighbor *bgp neighbor* family inet flow]** hierarchy level.

  To configure BGP to use VRF.inet.0 table to resolve VRF flow specification routes, include **secondary-independent-resolution** statement at the **[edit protocols bgp neighbor family flow]** hierarchy level.

  [See legacy-redirect-ip-action.]

  [See Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic.]

- **Support for 64 BGP add-path routes (MX Series)**—Starting in Junos OS Release 18.4R1, support is extended to 64 BGP add-path routes. Currently Junos OS supports six add-path routes and BGP can advertise up to 20 add-path routes through policy configuration. If you enable advertisement of multiple paths to a destination or if you increase the add-path prefix policy send count, BGP can now advertise up to 64 add-path routes.

  To advertise all add-paths, up to 64 add-paths or only equal-cost paths, include the **path-selection-mode** statement at the **[edit protocols bgp group *group-name* family *name* addpath send]** hierarchy level. You cannot enable both **multipath** and **path-selection-mode** at the same time.

  To advertise a second best path as a backup path in addition to the multiple ECMP paths include the **include-backup-path *backup_path_name*** statement at the **[edit protocols bgp group *group-name* family *name* addpath send]]** hierarchy level.

  [See path-selection-mode.]

  [See include-backup-path.]

- **Support for BGP egress peer engineering (MX Series)**—Starting in Junos OS Release 18.4R1, BGP LS extensions are enhanced to export segment routing topology information to the controller. A centralized controller in a software-defined network (SDN) can program any egress peer policy at ingress border

routers or at hosts within the domain in a segment routing network. The egress router advertises SID labels for all its peers, and the controller advertises these SID labels to the ingress router. The SID label can be a node segment, or an adjacency segment, or a set segment label. Thus the ingress router can select these SID labels to transfer data packets to the egress peers. The path that the controller derives can override the network derived best path. This feature can also be used in an inter domain scenario.

To configure a peer node SID, include **egress-te-node-segment-label** at the **[edit protocols bgp group *group-name* neighbor *neighbor-name*]** hierarchy level.

To configure a peer adjacency SID, include **egress-te-adj-segment *adj-segment-name*** at the **[edit protocols bgp group *group-name* neighbor *neighbor-name*]** hierarchy level.

To create a peer set SID, include **egress-te-set-segment *set-segment-name* label *label-name*** at the **[edit protocols bgp]** hierarchy level.

[See egress-te-node-segment.]

[See egress-te-adj-segment.]

[See egress-te-set-segment.]

- **Support for IPv4 VPN unicast and IPv6 VPN unicast address families in BGP (MX Series)**—Starting in Junos OS Release 18.4R1, the following address families are supported to enable advertisement or reception, or both, of multiple paths to a destination to and from the same BGP peer, instead of advertising and receiving only the active path to and 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**)

  [See Understanding the Advertisement of Multiple Paths to a Single Destination in BGP.]

- **BGP add path support for eBGP (MX Series)**—Starting in Junos OS Release 18.4R1, add path receive is now supported for eBGP under the [edit logical-systems *logical-system-name* protocols bgp group *group-name* family *family*].

  [See Understanding BGP.]

Services Applications

- **Support for MPLS-IPv6 inline active flow monitoring (MX Series)**—Starting in Junos OS Release 18.4R1 on MX Series routers, you can perform inline flow monitoring for MPLS-IPv6 traffic. Both IPFIX and version 9 templates are supported. If you are running inline flow monitoring on a Lookup (LU) card, you must enable sideband mode to create MPLS-IPv6 flow records.

  [See Configuring Inline Active Flow Monitoring Using Routers, Switches or NFX250.]

- **MX Series Virtual Chassis NAT support on BNG (MX240, MX480, and MX960 routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 18.4R1, you can configure a two-member MX Series Virtual Chassis to use the Juniper broadband network gateway (BNG) with IPv4-to-IPv4 basic NAT, dynamic NAT, static destination NAT, dynamic NAT with port mapping, and stateful NAT64. A two-member MX

Series Virtual Chassis configuration supports a maximum of four MS-MPCs and four MS-MICs per Virtual Chassis.

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

- **MX Series Virtual Chassis DS-Lite support (MX240, MX480, and MX960 routers with MS-MPCs and MS-MICs)**—Starting in Junos OS Release 18.4R1, you can configure DS-Lite on a two-member MX Series Virtual Chassis. A two-member MX Series Virtual Chassis configuration supports a maximum of four MS-MPCs and four MS-MICs per Virtual Chassis.

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

Software Defined Networking (SDN)

- **New features supported on Junos Node Slicing (MX Series)**—Starting in Junos OS Release 18.4R1, Junos Node Slicing supports the following features:
  - Support for device family and release in Junos OS YANG modules. [See Understanding Junos OS YANG Modules.]
  - Support for adding user-defined YANG files that provide mappings between the XML path and the OpenConfig path for data streamed through the Junos Telemetry Interface. [See Configurable NETCONF Proxy for Junos Telemetry Interface.]
  - Support for multiple, smaller configuration YANG modules. [See Understanding the YANG Modules That Define the Junos OS Configuration.]
  - Support for bidirectional authentication (client and server authentication) for gRPC for Junos Telemetry Interface. [See gRPC Services for Junos Telemetry Interface.]
  - Junos events sensor for the Junos Telemetry Interface. [See Overview of the Junos Telemetry Interface.]
  - Input streaming for gRPC Network Management Interface. [See Understanding OpenConfig and gRPC on Junos Telemetry Interface.]
  - ON_CHANGE support for Junos Telemetry Interface. [See Understanding OpenConfig and gRPC on Junos Telemetry Interface.]
  - Enhanced TACACS+ behavior to support the management interface in a non-default virtual routing and forwarding (VRF) instance. [See Management Interface in a Non-Default Instance.]

- TACACS+ authorization for operational commands using regular expressions. [See Using Regular Expressions on a RADIUS or TACACS+ Server to Allow or Deny Access to Commands.]
- Enhanced support for the nondefault management instance mgmt_junos. [See Management Interface in a Non-Default Instance.]

Subscriber Management and Services

> **NOTE:** Subscriber management is not ready for deployment in Junos OS Release 18.4R1. You can use this release for testing and qualification, but we recommend you wait for a later 18.4 maintenance or service release for deployment.

- **Limit subscriber sessions per user and access profile (MX Series)**—Starting in Junos OS Release 18.4R1, you can configure a limit on the number of sessions that can be active for a given username in an access profile.

  The **show network-access aaa statistics session-limit-per-username** command displays the number of active sessions and of blocked requests for usernames in each access profile. The **clear network-access aaa statistics session-limit-per-username** command enables you to clear blocked requests for debugging subscriber session limits.

  [See Understanding Session Options for Subscriber Access.]

- **New BBE statistics collection and management process (MX Series)**—Starting in Junos OS Release 18.4R1, the BBE statistics collection and management process, bbe-statsd, is introduced to take advantage of high-performance Routing Engines to increase the frequency of statistics collection and improve statistics processing in highly scaled environments. The **bbe-stats-service** option has been added to the **restart** command for restarting this statistics process.

  To collect subscriber and service statistics, you now must enable the **actual-transit-statistics** statement. If you do not configure this statement, subscriber statistics are not collected; the **show subscribers accounting-statistics** command displays a value of zero for subscriber statistics; and the subscriber statistics are reported to RADIUS with values of zero.

  [See Enabling the Reporting of Accurate Subscriber Accounting Statistics to the CLI.]

- **Subscriber secure policy information not revealed in core file dumps (MX Series)**—Starting in Junos OS Release 18.4R1, subscriber secure policy (SSP) information that might identify subscribers or mediation devices is automatically encrypted when the authd, bbe-smgd, or dfcd process generates core error files. Unauthorized persons examining the error files are unable to view the SSP information. The SSP information that might be present in the core error file includes the source and destination IP address for the mediation device, device ports, and intercept ID. No configuration is required or possible.

  [See Subscriber Secure Policy Overview.]

- **Increased number of IP addresses in DHCPv4 server groups (MX Series)**—Starting in Junos OS Release 18.4R1, DHCPv4 server groups support up to 32 active server IP addresses. In earlier releases, only 5 servers are supported.

  [See Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups.]

- **Address allocation method determines behavior when address pool is deleted or drained (MX Series)**—Starting in Junos OS Release 18.4R1, additional checking is performed to determine the subsequent behavior when authd notifies the DHCP process that an address pool is deleted or being drained:

  - When addresses are allocated on demand, the family with the address in that pool is logged out immediately when the pool is deleted, or logged out gracefully by the draining process when a DHCP renew or rebind message is received.
  - When the addresses are preallocated, the addresses for both families are deleted immediately when the pool is deleted, or deleted gracefully by the draining process when a DHCP renew or rebind message is received.

  [See Single-Session DHCP Dual-Stack Overview and Configuring DHCP Local Address Pool Rapid Drain.]

- **Enhanced support for forwarding ACKs from trusted servers (MX Series)**—Starting in Junos OS Release 18.4R1, the **allow-server-change** option of the **active-server-group** statement enables the DHCPv4 relay agent to forward ACKs to DHCP information request (DHCPINFORM) messages from any server in the active server group to the client. In earlier releases, only ACKs to DHCP request (renew or rebind) messages can be forwarded from trusted servers.

  [See Configuring Active Server Groups to Apply a Common DHCP Relay Agent Configuration to Named Server Groups.]

- **Support for DHCPv6 NotOnLink status code (MX Series)**—Starting in Junos OS Release 18.4R1, the DHCPv6 server can return to the client a status code of NotOnLink in the Reply PDU IA field during reauthentication when the subscriber IP or IPv6 address changes. This code means that at least one address in the client's request IA is not appropriate for the client's connection link. In earlier releases, only a NoAddrsAvail or NoPrefixAvail status code can be returned when there is an issue with requested addresses.

  [See RADIUS Reauthentication As an Alternative to RADIUS CoA for DHCP Subscribers.]

- **Reassign IPv4 address to a new subscriber (MX Series)**—Starting in Junos OS Release 18.4R1, you can enable a new subscriber to be reassigned an IPv4 address that is currently assigned to an existing subscriber by including the **reassign-on-match** option with the **address-protection** statement. The new subscriber request is rejected, but the existing subscriber is disconnected. The address is assigned to the new subscriber when it renegotiates the session

  [See Configuring Duplicate IPv4 Address Protection for AAA.]

- **New predefined variables and RADIUS VSAs for interface and set targeted distribution (MX Series)**—Starting in Junos OS Release 18.4R1, when you target an interface or an interface set for

distribution on aggregated Ethernet member links, you can use a Juniper Networks predefined variable to source the weight value from the RADIUS Access-Accept message on a per-subscriber basis, or from Diameter AVPs during NASREQ processing:

- $junos-interface-target-weight corresponds to Juniper Networks VSA 26-214, Interface-Targeting-Weight.
- $junos-interface-set-target-weight corresponds to Juniper Networks VSA 26-213, Interface-Set-Targeting-Weight.

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

- **Support for exporting BNG sensor data to an IPFIX collector (MX Series)**—Starting in Junos OS Release 18.4R1, the input-jti-ipfix plug-in collects a limited set of sensor data from the local BNG Junos Telemetry Interface and translates it to the appropriate IPFIX records for export to an IPFIX collector.

  [See Telemetry Data Collection on the IPFIX Mediator for Export to an IPFIX Collector.]

- **Detection and autogeneration of logical interface sets representing logical access nodes (MX Series)**—Starting in Junos OS Release 18.4R1, you can configure the router to parse the ANCP Access-Aggregation-Circuit-ID-ASCII attribute (TLV 0x0003). When the TLV string begins with a **#** character, the entire string is a backhaul line identifier. The portion of the string after the **#** delimiter represents a logical intermediate node (DPU-C or PON tree) in the access network to which the subscriber is attached. This portion is used to set the value of the $junos-aggregation-interface-set-name variable, and is used as the name of a CoS Level 2 interface set that groups subscribers. Enable parsing with the **hierarchical-access-network-detection** option of the **access-line** statement.

  [See Detection of Backhaul Line Identifiers and Autogeneration of Intermediate Node Interface Sets.]

- **BGP support over dynamic PPPoE interfaces (MX Series)**—Starting in Junos OS Release 18.4R1, BGP is supported over dynamic PPPoE interfaces. PPPoE subscriber clients correspond to BGP neighbors, so you configure the PPPoE subscriber client IP addresses as the BGP neighbor addresses with the **[edit protocols bgp group *name* neighbor]** stanza.

  You must enable routing services in both the PPPoE subscriber dynamic profile and the dynamic profile for the underlying VLAN interface with the new **routing-service** statement. This statement replaces the deprecated **routing-services** statement.

  You can also selectively enable or disable routing services per subscriber through RADIUS by using the new $junos-routing-services predefined variable. The action is determined by the value of the new Routing-Services VSA (26-212) returned in the RADIUS Access-Accept message.

  [See Junos OS Enhanced Subscriber Management.]

- **Support for Layer 2 services provisioning on the services side of pseudowire service logical interface anchored on redundant logical tunnel interface (MX Series with MPC and MIC)**—Starting in Junos OS Release 18.4R1, Layer 2 services provisioning such as bridge and VPLS, is supported on the services side of the pseudowire service logical interface anchored to redundant logical tunnel interface. With this support, the chassis-wide scaling numbers available for the physical interfaces over redundant logical tunnels is extended to pseudowire service interfaces anchored over redundant logical tunnel interfaces.

[See Layer 2 Services on Pseudowire Service Interface Overview.]

- **Support of single-hop BFD sessions for pseudowire redundant logical interfaces (MX Series)**—Junos OS supports inline distribution of single-hop Bidirectional Forwarding Detection [protocol] (BFD) sessions for pseudowire subscriber logical tunnel interfaces by default, as these interfaces are anchored on a single Flexible PIC Concentrator (FPC). With pseudowire redundant logical interfaces, the member logical tunnel interfaces can be hosted on different linecards. As a result, single-hop BFD sessions are operated in a centralized mode because the distribution address is not available for these logical interfaces.

  Starting in Junos OS Release 18.4R1, the support for inline distribution of single-hop BFD sessions is extended to pseudowire subscriber over redundant logical tunnel interfaces, thereby improving the scaling (number of sessions) and performance (detection time) of single-hop BFD sessions.

  [See Anchor Redundancy Pseudowire Subscriber Logical Interfaces Overview.]

- **ARP enhancements for subscriber management (MX Series)**—Starting in Junos OS Release 18.4R1, the following ARP enhancements are supported only for framed routes on dynamic VLANs:

  - Dynamic layer 2 MAC address resolution works for network (non-host) IPv4 framed routes. The non-host framed route is coupled with the dynamic Layer 2 address associated with a host route.
  - You can enable the router to compare the source MAC address received in a gratuitous ARP request or reply packet with the value in the ARP cache. The router updates the cache with the received MAC address if it determines this address is different from the cache entry.
  - You can enable dynamic ARP to resolve the MAC address for IPv4 framed host (32-bit) routes. By default, the framed route is permanently associated with the source MAC address received in the packet that triggered creation of the dynamic VLAN.

    [See Junos OS Enhanced Subscriber Management.]

System Management

- **Secure copy (scp) support on Junos OS CLI with the ”source address” and ”routing instance” options (MX240, MX480, MX960, MX2010, MX2020, and vMX)**— Starting in Junos OS Release 18.4R1, MX Series routers support the **scp** command from the CLI, along with two additional options: **source address** and **routing instance**. The **source address** option specifies the local address to use in originating the connection and **routing instance** option specifies the name of routing instance for the scp session. These two options are also added in the following CLI commands where the scp URL is supported: **file copy**, **file archive**, **save**, **show|save**, **show|compare**, **load merge**, **load override**, **load patch**, **load replace**, **load set**, and **load update**. The functionality of these commands remains the same with the **source address** and **routing instance** options added.

  NOTE: The scp command is available under operational mode and configuration mode.

  [See scp , file copy, file archive, load, and save.]

Timing and Synchronization

- **Synchronous Ethernet support for enhanced Switch Control Board (MX240, MX480, and MX960)**—Starting in Junos OS Release 18.4R1, MX Series routers with the enhanced Switch Control Board (SCBE3-MX) support synchronous Ethernet. Synchronous Ethernet is a physical layer technology that functions regardless of the network load and supports hop-by-hop frequency transfer. This enables you to deliver synchronization services that meet the requirements of modern-day mobile network, and future Long Term Evolution (LTE)–based infrastructures.

  [See Synchronous Ethernet Overview.]

VPN

- **Support to control traceroute over Layer 3 VPN (MX Series)**—Starting in Junos OS Release 18.4R1, in a Layer 3 VPN topology with **vrf-table-label** configured and multiple customer edge (CE) routers configured in the same VPN routing and forwarding (VRF) routing instance, when traceroute is performed to a remote provider edge (PE) router for a CE-facing network, the ICMP time exceeded packet determines the correct IP address as the source address.

  To control the traceroute over Layer 3 VPN topology with vrf-table-label configured and multiple CE routers configured in the same VRF, you can configure **allow-l3vpn-traceroute-src-select** at the**[edit system]** hierarchy level that determines the correct IP source address by reviewing the destination routing instance and destination IP address.

  [See allow-l3vpn-traceroute-src-select.]

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 main release and the maintenance releases for the MX Series routers.

Release 18.4R2-S6 Changes in Behavior and Syntax

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.

  [See interface-transmit-statistics.]

Release 18.4R2-S3 Changes in Behavior and Syntax

Platform and Infrastructure

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

Release 18.4R2-S1 Changes in Behavior and Syntax

Software Defined Networking (SDN)

- **Increase in the maximum value of delegation-cleanup-timeout (MX Series)**—You can now configure a maximum of *2147483647* seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

  With the increase in maximum value of **delegation-cleanup-timeout** from *600* to *2147483647* seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

  [See delegation-cleanup-timeout.]

Release 18.4R2 Changes in Behavior and Syntax

EVPN

- **Support for an VNI of zero**—Starting with Release 18.4R2, 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 18.4R2, 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.

Interfaces and Chassis

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (MX Series)**—In Junos OS Release 18.4R2, 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 before an interface changes state from down to up. In earlier Junos OS releases, the LACP hold-up the information for all interfaces was in a single **<lacp-hold-up-information>** XML tag. Now, the hold-up information for each interface is displayed in a separate **<lacp-hold-up-information>** XML tag.
- **Support for MAP-E de-encapsulation and decapsulation on inline service interfaces (MX2010)**—In Junos OS Releases 18.2R3, 18.3R2, and 18.4R2, MX2010 routers support encapsulation and de-encapsulation 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)
- **IRB not supported on pseudowire subscriber (PS) logical interface in bridge-domain (MX Series)**—In Junos OS Releases 17.4R3, 18.1R4, 18.2R3, 18.3R2, and 18.4R2, Integrated routing and bridging (IRB) is not supported on Pseudowire Subscriber (PS) Logical Interface. Thus you cannot add an IRB to bridge domain with a pseudowire subscriber interface–that is, you cannot configure IRB and the pseudowire subscriber interface in the same bridge domain.

  Note that adding IRB to a bridge domain having a pseudowire subscriber logical interface causes kernel crash and continuous reboot of the router until the configuration is rolled back.

  > **NOTE:** IRB is not supported on pseudowire subscriber interfaces only in bridge domain.

  [See bridge-domain.]
- In MX204 routers, error messages are logged when **vlan-tagging** for a trunk interface that is not configured. These error messages were previously logged with the 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 an informational error message (instead of an critical error message).

Junos OS XML API and Scripting

- **Root XML tag change for show rsvp pop-and-forward | display xml command (MX480)—**We've changed the root XML tag for the show rsvp pop-and-forward | display xml command to rsvp-pop-and-fwd-information to make it consistent with the XML tag convention. In earlier releases, the command output displays rsvp-pop-and-fwd-info XML tag. Update the scripts with the rsvp-pop-and-fwd-info XML tag to reflect the new rsvp-pop-and-fwd-information XML tag.

  [See Junos XML API Explorer - Operational Tags.]

Operation, Administration, and Maintenance (OAM)

- **Performance monitoring history data is lost when change in number of supported history records is detected (MX Series)**—In Junos OS Release 18.4R2, when Ethernet connectivity fault management (Ethernet CFM) starts, it detects the number of history records supported by the existing performance monitoring history database if there is any change from the number of history records supported (that is, 12) in Releases 18.4R2, then the existing performance monitoring history database is cleared and all performance monitoring sessions are restarted with mi-index 1.

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.]

Services Applications

- **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. If 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 NAT pool address.

- **Support for host-generated traffic on a GRE-over-GRE tunnel (MX Series)**—In Junos OS Release 18.4R2, you can send host-generated traffic on a GRE-over-GRE tunnel. However, when the path maximum transmission unit (path MTU) is updated for the outer GRE tunnel, MTU for the inner GRE tunnel is not corrected.

- **Deprecated IPsec manual security association option (MX Series)**—In Junos Release 18.4R2 and later releases, the option **hmac-sha2-256** under the **services ipsec-vpn rule *rule-name* term *term-name* then**

manual direction (bidirectional | inbound | outbound) authentication algorithm statement is deprecated. Use the **hmac-sha-256-128** option instead.

- **Change in error message displayed while fragmenting or de-fragmenting IPv6 GRE tunnel interface (MX Series routers)**—In Junos OS Release 18.4R2, on an 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 releases before Junos OS 18.4R2 release, the following message is displayed:

  dcd_config_ifl_tunnel: Fragmentation for V6 tunnels is not supported

Subscriber Management and Services

- **Out-of-address SNMP trap requires thresholds to be configured (MX Series)**—Starting in Junos OS Release 18.4R2, 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 level. 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.

- **Changing attributes of physical interface with active subscribers (MX Series)**—Starting in Junos OS Release 18.4R2, the commit check fails when you change any attribute of the physical interface, such as the MTU, when subscribers are active. This affects only aggregated Ethernet physical interfaces with targeted distribution configured. In earlier releases, the commit check does not fail and the attribute change brings down the physical interface and all subscribers using that interface.

- **Subscribers allowed to log in with bad framed route (MX Series)**—Starting in Junos OS Release 18.4R2, users are allowed to log in if the framed route received from RADIUS is bad–for example, if the format is incorrect. In earlier releases, the subscriber is not allowed to log in. For customers that use multiple framed routes, the new behavior enables the subscriber to have partial access to the network using the routes that are accepted instead of not being allowed any access.

- **ICMP error message rate limit increased (MX Series)**—Starting in Junos OS Release 18.4R2, 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.

Release 18.4R1 Changes in Behavior and Syntax

General Routing

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

Interfaces and Chassis

- **New option to configure IP address to be used when the Routing Engine is the current master**—Starting in Junos OS Release 18.4R1, 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 allows 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 releases, same IP address would be applied on master and backup Routing Engines resulting in configuration issues.
- **TLV status for Layer 2 protocols (MX480)**—Starting in Junos OS Release 18.4R1, the output fields **Next-hop** and **vpls-status** are displayed in the **show interfaces *interface name* detail** command, only for Layer 2 protocols on MX480 routers.
- **Enhanced AC PEM in high-line power configuration supplies 2400 W power (MX240)**—Starting in Junos OS Release 18.4R1, on MX240 routers, the enhanced AC PEM in high-line power configuration provides a power output of 2400 W. On Junos OS versions prior to 18.4R1, the PEM provided only 2050 W of power output.

  [See show chassis power.]
- **Support for creating layer 2 logical interface independently (MX Series)**—In Junos OS Releases 18.4R1, 18.4R2, and later, MX Series routers support creating layer 2 logical interface independent of layer 2 routing instance type. That is, you can configure and commit the layer 2 logical interfaces separately and add the interface to bridge-domain or Ethernet VPN (EVPN) routing instance separately. Note that the layer 2 logical interfaces works fine only when the interface is added to bridge domain or EVPN routing instance.

In the earlier Junos OS releases, when an layer 2 logical interface configuration (units with encapsulation vlan-bridge configuration) is used, then the logical interface must be added as part of a bridge-domain or EVPN routing instance for the commit to succeed.

- **Error message displayed due to configuration changes in live system**—Starting in Junos OS Release 18.4R1, on MX Series routers with the RE-S-X6-64G and RE-MX2K-X8-64G Routing Engines, when you change the router configuration on a live system, or when you delete an interface that has active traffic, the message **select: protocol failure in circuit setup** is randomly displayed. However, there is no known functional impact.

MPLS

- 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 18.4R1, when you configure zero as the bandwidth, **0** is applied as the RSVP bandwidth.

  [See bandwidth (Protocols RSVP).]

- Starting in Junos OS Release 18.4R1, the remote procedure call (RPC) protocol XML tag for **mpls-label-value** is renamed as **mpls-history-label-value**, **mpls-usage-label-value**, and **mpls-label-id-value** depending on the context of command usage.

- **Change in command syntax**—Starting in Junos OS Release 18.4R1, the **show ldp database label-requests** command name is changed to **show ldp database-label-requests** with no change to command functionality.

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

Network Management and Monitoring

- **SSHD process authentication logs timestamp (MX Series)**—Starting in Junos OS Release 18.4R1, the SSHD process authentication logs use only the time zone defined in the system time zone. In the earlier releases, the SSHD process authentication logs sometimes used the system time zone and the UTC time zone.

  [See Overview of Junos OS System Log Messages.]

- **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 18.4R1, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, the server must not return an RPC reply that encloses both an **<rpc-error>** element and an **<ok/>** element. If the operation is successful, but the server reply would enclose one or more **<rpc-error>** elements of severity warning in addition to the **<ok/>** element, then 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 encloses both an **<rpc-error>** element of severity warning and an **<ok/>** element.

- **Change in severity level of XQSS errors (MX Series)**—Starting in Junos OS Release 18.4R1, 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 longer cause the entire FPC to go offline by default. Instead, these errors cause the affected Packet Forwarding Engine (PFE) to be disabled, because **disable-pfe** is the default action associated with Major errors on MX Series routers.

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

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

  [See show chassis fpc errors.]

Routing Protocols

- **BGP PIC determines MPLS fast reroute (FRR) using BPG multipath**—Starting in Junos OS Release 18.4R1, when you configure BGP Prefix-Independent Convergence (PIC) with the **protect-core** statement, a forwarding route with an MPLS fast reroute (FRR) next hop is created using BGP multipath.

  In earlier releases, when the BGP PIC feature is configured, a backup path is determined using protocol-independent load balancing multipath and installed in the forwarding table as an active path, which might cause routing loops.

  We recommend that you update scripts that count active routes because BGP multipath contributors are also counted and the active route count goes up. We have also modified the output of the **show route** command to reflect this behavior change.

  [See Configuring BGP Prefix Independent Convergence for Inet.]

Security

- **Syslog updated when configuring XPN cipher suite on a non-xpn supported interface (MX Series)**—In Junos OS Release 18.4R1, on MX Series Routers, if you attempt to configure XPN cipher suite (gcm-aes-xpn-128 or gcm-aes-xpn-256) for a connectivity association and attach the connectivity association to an interface on the PIC that does not support XPN cipher suite, then during runtime, a syslog is logged as below (and default non-xpn cipher suite is used):

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

Software Defined Networking (SDN)

- **Installation or upgrade using remotely located installation package (MX480, MX960, MX2010, MX2020, MX2008)**—While performing Junos 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

- **ZTP is supported on MX PPC platforms (MX Series)**—As of Junos OS Release 18.4R1, 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

- **Flat-file service accounting support ends (MX Series)**—Starting in Junos OS Release 18.4R1, flat-file service accounting to a local file is no longer supported. If included in a configuration, it is ignored.

  [See Flat-File Accounting Overview.]

SEE ALSO

New and Changed Features | 90

Known Behavior | 119

Known Issues | 123

Resolved Issues | 140

Documentation Updates | 172

Migration, Upgrade, and Downgrade Instructions | 173

Product Compatibility | 180

Known Behavior

IN THIS SECTION

This section contains the known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 for MX Series routers.

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

Forwarding and Sampling

- LTS subscriber statistics are reported to RADIUS. PR1383354
- In Junos OS Release 18.4R1 and Release 18.3R2, if iPv4 prefix is added to a prefix list referred to by IPv6 firewall filter, then the log message **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized** is not seen in this particular release. PR1395923

General Routing

- When a route or a next hop has been created by the application, we assume that it can propagate to the rest of the system. KRT asynchronously picks up this state for propagation. There is no reverse indication to the application if there was an error in propagating the state. The system is supposed to eventually reconcile. So, if SPRING-TE produces a <route, NH> pair that looks legal from the application's standpoint, but the KRT is not able to download it to the kernel (because kernel rejected the next hop), the <route, NH> pair gets stuck in the rpd. In the meantime, the previous version of the route (L-IS-IS in this case) that was downloaded still lingers in the kernel and Packet Forwarding Engine. PR1253778
- CFM is not supported for an L2-over-GRE tunnel. CCM can pass through as transit traffic through GRE interfaces transparently using the data path. Link trace functionality uses MAC-learning and re-injecting LTM on the GRE interfaces in case the bridge is configured with CFM. PR1275833
- An underflow error is seen during FPC cold boot and initial traffic start cases. But these errors are limited and should not appear after traffic is stabilized. PR1306280
- Support for enterprise profile is provided only for 10-Gigabit Ethernet interfaces. Use of 40-Gigabit Ethernet and 100-Gigabit Ethernet interfaces might result in a phase alignment issue. PR1310048
- When cmerror disables the Packet Forwarding Engine, it does not power off the ea and hmc chips. Temperature monitoring continues on the HMC and other devices, and the system can take proper actions, such as increasing the fan speed or shutting 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
- The Routing Engine boots from the secondary disk when you:

  Press the reset button on the RCB front panel while the 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.

  Reboot, and the system hangs before Junos OS is up. PR1344342

- The MIC-MACSEC-20G MIC supports 10-Gbps speed through the **set chassis fpc x pic y pic-mode 10G** configuration applied to both the PICs in that MIC. Any other PIC mode configuration should be removed before you apply the 10G PIC mode configuration. PR1374680
- In Junos OS most daemons (Junos OS processes) underwent architectural change in transition from Junos OS Release 14.1X53 to Junos OS Release 17.x (4 years) and many new features were added. These changes caused an increase in memory footprint in 17.X compared to Release 14.1X53. Unless we see system instability or any adverse performance impact, or a daemon crash due to low memory, this increased memory footprint should not be an issue, and functionality should work fine. The increased memory footprint is a Junos OS property not specific to QFX5000. PR1390226
- IDS aggregate configuration statement is not considered for the installation of the IDS dynamic filter. PR1395316
- Junos OS does not perform the VLAN ID check at the egress; the VLAN ID check is performed only at ingress. PR1403730

Interfaces and Chassis

- During JDM installation, each JDM instance generates pseudorandom MAC addresses to be used for JDM's own management interface and for the associated GNFs' management interfaces. At the time of creation of GNFs, each GNF instance generates pseudorandom MAC addresses to be used as the chassis MAC address pool for the forwarding interfaces of that GNF. Once generated, JDM and GNF MAC addresses are persistent, and are 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.
- The two SFP+ ports on the Routing Control Board (RCB) of an MX2008 router have two port LEDs each—one **Link Status** LED and one **Link Activity** LED per port. On an MX2008 router, which is connected to an external x86 server in a Junos node slicing setup, behavior of these LEDs with regard to Junos Node Slicing configuration is as follows:
  - The **Link Status** LEDs and **Link Activity** LEDs on both the ports are unlit when Junos node slicing is disabled or not configured.
  - When you have configured **network-slices** on the router (also called base system or BSYS) but have not configured guest network functions (GNFs) on the server, the **Link Status** LED on each port turns green (steady glow). In this case, the **Link Activity** LED on each port is unlit.

  - When you have configured Junos node slicing (including GNFs), the **Link Activity** LED on each port is amber (blinking), while the **Link Status** LED on each port remains green (steady glow).

- **Error thrown when router configuration is updated on live system**—In Junos OS Release 18.4R1, on MX Series routers with the RE-S-X6-64G and RE-MX2K-X8-64G Routing Engines, when you change the router configuration on a live system, or when you delete an interface that has active traffic, the message **select: protocol failure in circuit setup** is randomly displayed. However, there is no known functional impact.

Platform and Infrastructure

- On all devices running Junos OS, execution of Python scripts through enhanced automation does not work on veriexec images. PR1334425
- A few transient FI Cell underflow errors are normal during unified ISSU, but they should not persist. PR1353904
- On QFX10000 switches configured as type-5 route peers, when only peer 1 advertises routes, that peer might not install the de-encapsulated next hop (NH) route. As a result, type-5 encapsulated traffic sent by peer 2 is dropped until peer 2 advertises any type-5 route. As a workaround, configure a static route pointing to discard on peer 2 and advertise that route as a type-5 route to peer 1. PR1386423
- In some cases, PS interfaces over RLT might be shown as up but they might not pass traffic. Log messages reporting ASIC errors and a chassis alarm reporting hard FPC errors may also be seen. PR1400269

Routing Protocols

- When multiple adjacencies are coming up or flapping, some routes may not have remote-lfa backup next hops. They will appear only after the next SPF trigger, either manually or as a result of a network event. PR1389392

Subscriber Management and Services

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

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

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.4R2 for MX Series routers.

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

Class of Service (CoS)

- Configuration of the hidden configuration statement **rate-limit-burst** in the class-of-service 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 setting (Internal attributes for an interface) was found to be NULL. Interface settings are usually stored in a memory location and the pointer to it became NULL because cosd process did not check for the NULL values and resulted in segmentation fault. Channelized interface setting was found to be NULL for channelized interfaces, but the CoS code handling the configuration **rate-limit-burst** configuration in the Packet Forwarding Engine de-referenced the setting without performing NULL check, thus resulting in core files. PR1425667

EVPN

- The issue is applicable to MAC-in-MAC PNN-EVPN and does not affect any other scenario. When the provider backbone bridging (PBB) EVPN configuration is reloaded on MX Series routers, error logs are seen while deleting interfaces related to backbone bridge component. These errors do not result in any functional issues. PR1323275
- Type 2 EVPN routes are missing after the EVPN protocol is deactivated and then reactivated. PR1362598
- Ping overlay - RPC error (illegal option ? X?). PR1373025
- When EVPN is configured with CoS-based forwarding (CBF), traffic might be lost for the CBF services. PR1374211
- Replace the multihome advertisement proxy bit from L2_info community to ARP/ND extended community. The default value is 0x4. PR1408055
- In an EVPN-VXLAN scenario with scaled bridge domains configured (for example, with 4000 bridge domains), if the core-facing link on the VXLAN tunnel endpoint (VTEP) comes up (Down >> Up), the traffic received from the customer edge might be dropped by the VTEP for a period of time before it becomes normal. PR1408840
- The core-isolation feature does not work after setting and then deleting the **no-core-isolation** statement on MX Series. The feature can be enabled back after restarting rpd. PR1442973

Forwarding and Sampling

- The **skip-service** configuration does not work with IPv6 NDP negotiation or ping. PR1074853
- Heap memory leaks occur on the DPC when the flow specification route is changed. PR1305977

- This PR is to fix some hints for the CLI commands to avoid confusion. With the fix, it should be as follows: **{master}[edit] user@host-re1# set firewall flexible-match source-ipv6-match bit-length ?**. PR1389103
- On Junos fusion, ingress policing on SD is broken (MX+QFX: Ingress policing on AD and SD) the **set interfaces layer2-policer input-policer** *policer-name* command is not supported in this release. PR1395217
- In Junos OS Release 18.4R1 and Junos OS Release 18.3R2, if IPv4 prefix is added on a prefix list referred to buy an IPv6 firewall filter, then the log message **Prefix-List [Block-Host] in Filter [Protect_V6] not having any relevant prefixes , Match [from prefix-list Block-Host] might be optimized** will not be seen in Junos OS Release 18.4R2. PR1395923

General Routing

- If a Layer 3 interface is receiving a GRE-encapsulated packet and the interface has two filters attached in ingress as follows:
  - . **Family any** with action as mirror
  - . **Family inet** with action as decapsulate gre, then the expected behavior is that the mirrored copy must have the GRE headers as well. However, that is not working as expected (and a bug) due to the presence of the family inet filter. If you are interested in mirroring entire packet that came on the interface (that includes GRE header as well), then workaround is to deactivate/disable the "decapsulate gre action of filter.PR1090854
- The nexthop attribute in a framed route is not applicable anymore. Because subscriber's IP address is used as the next hop in all cases, there is no need to have an additional nexthop attribute for framed routes. PR1186046
- During a Routing Engine switchover (without NSR), the l2cpd process might report a slip (delay) of 1--10 seconds in its scheduled run, and a log message similar to the following might be displayed: **Aug 1 10:41:21 mx9601 l2cpd[32770]: JTASK_SCHED_SLIP: 8 sec scheduler slip, user: 0 sec 2180 usec, system: 0 sec, 2188 usec**. This delayed run has no functionality nor operational effect to any of the Layer 2 protocols controlled by l2cpd because STP task delegates transmit/receive BPDUs to a separate dedicated ppmd process, and the LLDP task's transmit/receive PDUs are dealt from the daemon itself but the advertisement interval is 30 seconds, with the hold timer for the neighbors' LLDP PDU being 120 seconds. Thus, the time to recover the few seconds of delay is plenty and enough to absorb the delay. PR1203977
- In a rare race condition, multiple interrupts are not handled properly on MX Series devices with MPC7E, MPC8E, or MPC9E and PTX Series devices with FPC3-PTX-U2 and FPC3-PTX-U3, which could lead to the generation of 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
- In a BGP or an MPLS scenario, if the next-hop type of the label route is indirect, then the following changing events related to the family mpls configuration of the next-hop interface might cause the route to be in dead state, and the route remains dead even when the family mpls configuration is again activated.

Deactivating and activating the interfaces family mpls configuration

Deleting and adding back the interface's family mpls configuration

Changing the maximum-labels setting for the next-hop interface

> **NOTE:** When a labeled route is resolved over an interface, that interface must have family mpls configured for the route to be successfully resolved. Otherwise the route does not get resolved. PR1242589

- Load balancing is uneven across aggregate Ethernet member links when the aggregated Ethernet bundle is part of an equal cost multi-path (ECMP) path. The member links needs to span 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. Please open a JTAC case to confirm**. PR1258970
- If a VM host snapshot is taken on the alternate disk and there is no further vmhost software image upgrade, the expectation is that if the current VM host image gets corrupted, the system will boot from the alternate disk so that the user can recover the primary disk to restore the state. However, if the host root file system is corrupted, the node is booting with the previous VM host software as against booting from the alternate disk. PR1281554
- This issue is noticed with the 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
- You cannot collect shmlog entries and statistics on MX5, MX10, or MX40 platforms. The code changes also include improvements that should prevent the generation of shmlogctl process core files due to a timing issue. PR1297818
- The **show dynamic-tunnels database summary** command would not show accurate tunnels summary during the time anchor Packet Forwarding Engine linecard is not in up state. Use below commands as a work around: **show dynamic-tunnels database** and **show dynamic-tunnels database terse**. PR1314763
- As a vendor does not use chained CNH, using the feature does not bring in a lot of gain because TCNH is based on an ingress rewrite premise. Without this feature things work just fine. PR1318984
- In JDM, (running on the secondary server) the jdmd daemon might generate core files if adding an image for the GNF is aborted by pressing CTRL-C. PR1321803
- With regard to FPC restarts or Virtual Chassis splits, the design of MX Series Virtual Chassis infra relies on the integrity of the TCP connections, and the reactions to failure situations might not be handled in a graceful way. TCP connection timeout because of jlock hog crossing the boundary value (5 seconds)

causes bad consequences in the MX Series Virtual Chassis. Currently, there are no other easy solutions that can reduce this jlock hog other than enabling the marker infra in the MX Series Virtual Chassis setup. Unfortunately, there is no immediate plan on enabling marker as it was causing a lot of issues in the MX Series Virtual Chassis when we tried to enable it. PR1332765

- The first packet pertaining to the J-Flow Packet Forwarding Engine sensor in UDP mode is missing after a line-card reboot. PR1344755
- With graceful Routing Engine switchover (GRES) enabled in a subscriber environment, if subscribers are logging in and logging out very quickly, the service sessions in the session database of the backup Routing Engine sessions might be leaked. If the problem is not detected for long enough, the backup Routing Engine might not be able to synchronize with the master Routing Engine and thus will not be ready for GRES. PR1346300
- During a unified ISSU that warrants host upgrade, if the router is configured with 8 million IPv4 or IPv6 routes or more, upgrade might fail, resulting in FPC restart. PR1348825
- On a next-generation Routing Engine (NG-RE), failure of the Hardware Random Number Generator (HWRNG) leaves 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 silent discarding of traffic that was destined to the MPC. The only way to recover from this situation is to restart the MPC. The issue is not seen if you use the corresponding CLI commands to take the MIC offline and then bring it online. PR1350103
- On all devices running Junos OS, licenses might not take effect after you have successfully committed a license key configuration. PR1350302
- The EX9253 switch does not support interface ranges for channelized interfaces. You need to configure the interfaces individually. PR1350635
- During stress conditions, error log messages regarding addition, modification, or deletion of routes might be incorrect. PR1350713
- If an aggregated Ethernet interface is configured with **link-protection backup-state down** , the AE operational state of the interface is still up even though the member interfaces configured under the aggregated interface are down. This issue is specific to the **link-protection backup-state down** configuration for the aggregated Ethernet interface. PR1354686
- The issue occurs only when you deactivate and then activate an aggregated Ethernet link, which means that the LAG interface is deleted from the system and created again. But then, the issue does not happen on deactivating and 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/Backup Link/Primary Weight/Backup Weight> Group 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 gets 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 be with unilist->indirect->hold to unilist->indirect->unicast state transition during the deactivation-activation

process. As of now, we have a workaround to enable the vty command to change the unilist hold behavior. Because the issue gets replicated very sporadically, getting to the actual fix is taking some time. PR1354786

- On MX platforms with MS-MPC or MS-MIC, if a large sum of similar packets (for example, thousands of packets) are received, and because of the flaw of the method to process these packets, data/management path was completely blocked and dead-locked. Eventually, traffic might be blocked. PR1358019
- The configurations of bridging routing instances having AE IFLS(6400IFLs) and IRB instances, all from a single FPC, the CPU utilization of the FPC stays at 100% for 4 minutes. The behavior from PFEMAN of FPC has the processing time spiked on IF IPCs and this seems to be the case of MPC7E from Junos OS Release 16.1R1(or even earlier). After 4 minutes, the CPU utilisation comes down and the FPC is normal. Therefore, this scale configuration on MPC7E takes a little more than 4 minutes to settle. PR1359286
- In rare circumstances, a faulty SFP installed in a MX104 might cause the AFEB to be offline. The Backup Routing Engine and Fan Tray will also be in alarm. PR1360426
- Syslog is updated when the user tries to configure XPN cipher over a non-XPN supported platform such as MIC-MACSEC-20G even though commit goes through. PR1367722
- When FPC is booting up (either during unified ISSU or router reboot or FPC restart), i2c timeout errors for the SFP transceiver can be noticed. These errors occur because the I2C action is not completed as the device was busy. After the line card is up and all the I2C transactions to the device are all right, so no periodic failure is observed. There is no functional impact and these errors can be ignored. PR1369382
- After successfully delegating a locally configured LSP to a PCE, the router still displays 0 as the "Delegated" counter value in 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 an MIC on MPC5E. PR1370337
- When the MIC-MACSEC-20G MIC is in offline state after Fake-Kats initiation, the MIC has to be brought up by issuing chassisd restart. Attempting to bring the MIC online using the CLI could cause the MIC to go into a hardware error state. PR1374532
- When CBF (CoS-based forwarding) is enabled, due to the indexed next hop installation issue in the kernel, the rpd process might crash upon route flap and LSP flap. PR1374558
- I/O session used for communicating between threads is freed due to FSM state transition. After freeing the memory, the fields of the I/O session are used for tracing, which leads to the generation of rpd core files. PR1374759
- If any of the log message continuous to pop in MPC console, it indicates the presence of a faulty SFP or SFP+ transceiver, which is causing an I2C transaction from the main board CPU. There is no software recovery available to recover from this situation. These logs also indicates potential I2C transaction failure with any of the 10 ports available with GMIC2 in PIC 0 resulting in unexpected behaviors such as link not coming up or MIC itself not booting up on restart. **I2C Failed device: group 0xa0 address 0x70Failed to enable PCA9548(0x70):grp(0xa0)->channel(0)mic_sfp_select_link:MIC(0/0) - Failed to**

enable PCA9548 channel, PCA9548 unit:0, channel ID: 0, SFP link: 0mic_sfp_id_read: Failed to select link 0 Only way to recover from these failures is to detect & replace faulty SFP/SFP+ plugged into the GMIC2 ports PR1375674

- Interface with Tri Rate Copper SFP(P/N:740-01311) in MIC 3D 20x 1GE(LAN)-E,SFP will stop forwarding traffic after ISSU upgrade. PR1379398
- In a subscriber scenario, if the **service-accounting-deferred** statement is configured for a dynamic profile, and there is multicast to a large number of destinations on the same physical port, then FPC errors might be seen. PR1380566
- In rare situations at heavy traffic loads, the input frame check sequence counter might get incremented. PR1383009
- Users can still issue the command **set vmhost...** command although **permissions system-control** is not configured on the system class. PR1383706
- Commit should not be allowed if we are trying to delete the **physical-cores** configuration statement. However, there is no functional impact. PR1384014
- In low-end 32-bit systems, rpd has a lower level of available memory. We need a log message to alert the user when the average memory usage or transient memory usage exceeds thresholds. PR1387465
- During the zero-touch provisioning (ZTP) process, the default route is being cleaned up by code. Due to this, if a static default route is configured in the initial configuration (configuration file downloaded from the file server for ZTP), the route will fail to work. This might lead to ZTP failure or a device access issue after ZTP. PR1387724
- On an MX Series enabled with enhanced subscriber management, if the filter service is enabled for each subscriber, and there is a large scale of Broadband Edge (BBE) subscribers (for example, 10000) logging in and out repeatedly, the FPC might crash due to this rare issue. PR1388120
- In cases of PS over rlt at high scale, removing and adding back a CoS configuration can cause the FPC to enter a hard error state. PR1388487
- The virtio throughput remains the same for multi-queue and single-queue deployments. PR1389338
- In a Junos fusion for provider edge (MX Series) scenario, all the FPCs might restart after the changes to the VLAN/encapsulation on the extended port are committed if the **per-interface-per-member-link ingress** parameter is configured for the sourced routing statistic by using the command **set protocols isis source-packet-routing sensor-based-stats per-interface-per-member-link ingress**. PR1392071
- MPC card/afeb/tfeb with Channelized OC MIC can crash with the generation of core files. PR1396538
- The Junos OS rpd daemon has facilities to attempt to trap certain classes of nonfatal bugs by continuing to run, but leaving a "soft" core file. Leaving a soft core file is intended to be nondisruptive to routing and forwarding. This PR implements a mechanism by which users can disable soft core files. PR1396935
- The interface link stays down when we deactivate and then activate the channelized xe- (10-Gigabit Ethernet) interface configured with speed 1 Gbps (when using QSA adaptor) on MX10008

(JNP10K-LC2101 MPC) with line rate traffic flowing. As a workaround, we need to take the MIC offline and then bring it online to recover the link; this is a known issue. PR1397202

- The CLI command **show system firmware** might provide an unexpected output on some MX Series platforms such as the MX104. The **current version** might be shown as ?? instead of the correct version number. PR1398022
- The **$junos-framed-route-ipv6-address-prefix** variable for programming IPv6 routes is permitted only under the **routing-options**->**rib**->**access** configuration. PR 1384523 changed the code to avoid the incorrect mixing of IPv4 and IPv6 framed routes in the same configuration and force the V6 framed routes to be parsed only if they were in their correct **routing-options**->**rib** ->**access** stanza. Additionally, runtime warnings for invalid configuration IPv6 framed routes configuration were added in PR 1388737. PR1401144
- In a BGP PIC scenario, If a route, R1, resolves on top of multipath-route R2, where R2 has primary and backup indirect next hops, it will be better if a backup leg is not used for the resolution of R1. There is no impact on any existing CLI commands. The backup path is never used when the primary path is available. PR1401322
- The **sample-frequency** data type is changed from "milliseconds" to "seconds." PR1402197
- After upgrading Junos OS to Junos OS Release 17.2 or later releases, the **chained-composite-next-hop ingress l3vpn extended-space** statement cannot be configured any longer on a logical system. PR1402390
- When you initiate the image installation on the base system of a setup with node slicing enabled, the session gets terminated unexpectedly. PR1402643
- 1G configuration mode is not a unified ISSU supported configuration on Summit MX 3RU router. If that configuration is present on the Summit MX 3RU box, then the user has to remove the same before attempting unified ISSU. Otherwise the 1-Gigabit Ethernet configurations will not behave as expected after unified ISSU and traffic loss can be expected. Currently there is no warning or error message alerting the customer about the issue. This is applicable to Summit MX 3RU platform only. PR1405527
- The rpd process might crash after a nonforwarding route (that is, a route to an indirect next-hop association is a non-forwarding indirect next-hop) that is received from multiple protocols is resolved again by using the non-forwarding path. PR1407408
- MX104 MIC (MIC-MACSEC-20GE) supports Extended Packet Numbering (XPN) mode on 1-Gigabit and 10-Gigabit Ethernet interfaces. PR1409457
- 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
- On MX Series devices, 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 the PCC if the PCEP session goes down and gets re-established within **delegation-cleanup-timeout** period. PR1415224

- With NETCONF the xmlns attribute is printed twice for the **rpc <get-arp-table-information>** call to the router. PR1417269
- Certain JNP10008-SF and JNP10016-SF manufactured between July 2018 and March 2019 might have incorrect core voltage settings. The issue can be corrected by reprogrammed the core voltage and updating the setting in nvram memory. PR1420864
- On MX Series devices, with 1xCOC12 or 4XCOC3 used, if channelized interfaces are configured, FPC CPU overuse might be seen. PR1420983
- On MX platform, issuing the **show forwarding-options load-balance ..** command might cause a Packet Forwarding Edge wedge after a certain number of attempts (fewer than 200 in test), if the **destination-address** statement of the command matches the default route with the "discard" action. This is because a defect code causing internal flow errors is involved in that scenario. PR1422464
- On the MX204 platform, the allocation of MAC address for the second PIC in the FPC might fall out of the MAC address pool, which might further cause MAC conflict in the network. PR1422679
- Added support for SFP-T with QSA adapter in MX10003. PR1422808
- The issue is limited to DB related to MAC-MOVE scenario. When dhcp-security is configured, if MAC moves happen for multiple IPv4 and IPv6 clients, the jdhcpd might consume 100% CPU and later crash. PR1425206
- On all Junos platforms running 64-bit mode rpd, the rpd will crash continuously if any protocol authentication (like MD5 authentication for BGP/ISIS/OSPF) is used along with master-password. PR1425231
- Whenever the **show snmp mib walk jnxMibs** command is issued, the following logs are seen in 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 has no impact on production with respect to the below KB: https://kb.juniper.net/InfoCenter/index?page=content&id=KB24394 The KB has three old PRs for SRX/M10i which has fix in 12.x version. However we are noticing these log messages in 18.x version in MX204. PR1425411
- On MX platforms with ppp configured, if ppp session age out, the session does not get deleted, and hence, you cannot log in to a new session. The ppp traffic might be dropped because of the duplicate-protection feature on interface. And the IP address of the ppp interface is not pingable. PR1428212
- More number of MACs/MAC-IPs can get learnt if the MAC and MAC-IP limits are configured in a particular sequence. An example is shown below:

  1. Learn 50 remote entries.

  2. Configure a MAC limit of 20 (remote entries remain intact, this works as expected).

3. Learn 50 local entries At this point, no local entries must be learned, as MAC limit is 20. However, all 50 local MAC entries are learned, causing the MAC count to be 100, which is incorrect. The same issue will be seen for the MAC-IP limit as well. PR1428572

- Some non-Juniper 40-gigabit SFP transceivers might use 100-gigabit QSFP28 marking in their EEPROM, indicating CDR bypass mode, which enables the use of 100G optics at 40-Gbps speeds. On some 40-Gbps line cards, Junos OS detects an incorrect pluggable QSFP28 transceiver of type 0x11 (17 decimal) inserted into a QSFP+ of type 0x0d in the cage and reports this error to syslog. PR1434183

Infrastructure

- When there is a high route churn or when there is a high rate of route updates being pushed to the kernel, the **show interface** command might show delay or not show all statistics due to route updates being prioritized over statistics messages. PR1250328
- Junos OS can hang while trying to acquire the SMP IPI lock during a reboot when it is running as a VM on Linux and QEMU hypervisor. PR1359339
- 32-bit Routing Engine memory exhaustion leads to kernel crash. PR1378313

Interfaces and Chassis

- Out-of-sequence packets are seen with LSQ interfaces. PR1258258
- 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 the upgrade. This is because of the old version of **/var/db/cfm.db**. PR1281073
- 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, and MPC9. The errors have an impact only for DWDM PICs, which does not affect on these MPCs. Hence this syslog message can be safely suppressed. PR1349277
- The error message **ppman_cfm_start_inline_adj: Failed to add Inline adj for CFM, pkt-len=0** is observed in some cases. But there is no functional impact. Sessions/adjacency would get programmed inline subsequently. PR1358236
- With **ppp-service traceoptions** configured as: **user@router> show configuration protocols ppp-service traceoptions file jtac-jpppd.log size 1g files 10; level all; flag all; filter {user {"subscriber@domain.com";} }**, it is expected to see only PPP negotiation events belonging to the subscriber defined in the filter section. However, in releases affected by this issue several stings of logs related to other (non interested) subscriber may be seen. PR1370994
- LFM sessions toward scaled peers might flap during the switchover phase of a unified ISSU. PR1377761
- As part of the EOAM programming, the LM counters are allocated. When an interface is deleted, the EOAM LM counters needs to be cleared. This is done as part of EOAM punt deletion. But there are

scenarios where the prog punt delete processing is received, the and the logical interface is deleted in ukern. In this case the EOAM NHs are cleared but the LM counters are not freed. This can cause memory leak in jnh. This issue is seen for a scaled configuration, with repeated addition and deletion of the interface configurations when EOAM configuration is present. PR1396540

- There might be a memory leak on **tarnsportd** when bulk SNMP polling occurs on large-scale logical interfaces and large number of traps are created because of an interface flapping. The memory leak could cause the transportd consuming high CPU for a prolonged period. PR1398967
- Static demux0 logical interfaces do not come up after a configuration change if the underlying interface is et- (that is, 100-Gigabit Ethernet). After the configuration change, the et interface gets flushed in order to reparse the configuration. During this DCD miss to create the dependency between demux0 logical interfaces and 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 has been already taken care of. This is day one issue. As a workaround, restart DCD (or reboot the entire Routing Engine). If the problem is not resolved, then use **commit full** instead of **commit** while committing new configuration. PR1401026
- On MX Series platforms, EX-SFP-1FE-LX SFP does not initialize with MIC-3D-20GE-SFP-E(EH). PR1405271
- When an unnumbered interface is binding to an interface that has more than one IP address and one of the IP addresses is deleted, the family **inet** of the unnumbered interface might get 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 addresses, hence avoiding the deletion of the family **inet** of the unnumbered interface. PR1412534
- If an aggregated Ethernet (ae-) interface has VRRP configuration, in the following use cases, member llogical interfaces are not created after the member physical interface comes up and the ae- interface is in down state:

  1. FPC restart (**request chassis fpc restart slot <>**)
  2. Chassis-control restart (**restart chassis-control**)
  3. Reboot both Routing Engine (**request system reboot both-routing-engines**)

  So before performing above operations, it is advisable to remove vrrp configuration from aggregated Interface(ae). PR1429045

Layer 2 Features

- In 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. It would cause the traffic loss for one CE site to peer CE site. If **connectivity-type permanent** is configured, this issue will not be observed as the instance will remain in 'UP' state. PR1415522

Layer 2 Ethernet Services

- In MC-LAG with **force-up** scenario, an LACP PDU loop might be seen when both MC-LAG nodes and the access device and use the same admin key. PR1379022
- On MX Series devices, if a static demux interface is configured over underlying, after subscriber logout, the accounting statistics are not cleared. PR1383265

MPLS

- If the primary link goes down immediately after bypass (for example, FPC containing both primary and bypass or, both primary and bypass FPCs go down simultaneously) such that the 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 are left with stale LSP state until refresh timeout. This condition does not result in any traffic loss. PR1242558
- With nonstop active routing (NSR), when the rpd restarts on the master Routing Engine, the rpd on the backup Routing Engine might also restart. PR1282369
- In case of CSPF-disabled LSPs, if the primary path ERO is changed to an unreachable strict hop, sometimes the primary path stays up with the old ERO. The LSP does not switch to standby secondary. PR1284138
- For an SR-TE path with "0" explicit NULL as the innermost label, the SR-TE path does not get installed with the label "0". PR1287354
- When you run a traceroute to a remote host for an MPLS LSP by using the **traceroute mpls bgp** command, in very rare cases, the mplsoam process might hold the stale BGP instance handle in the query to the rpd process to get the information for the forwarding equivalence class. As a result, rpd crash might occur because of the invalid instance. Traffic flow might be impacted until rpd comes back up. PR1399484
- On devices running Junos OS, with transit chaining mode enabled, if RSVP link/node protection is enabled and **sensor-based-stats** is used, a single-hop bypass label-switched path (LSP) next hop might not be installed in forwarding information base (FIB) even it is in the routing information base (RIB). Hence the single-hop bypass LSP might fail to forward traffic when needed. PR1401152
- With NSR enabled, when master rpd is restarted, occasionally, out-of-order add and delete messages can arrive on the backup Routing Engine causing label assignment that can result in rpd crash on the backup Routing Engine. PR1401813

- When make-before-break (MBB) new instance signaling experiences an error and before the retry is completed, other triggers such as automatic bandwidth adjustment timer expiration need to be blocked until MBB finishes. After the MBB finishes instance switching, blocked trigger needs to be scheduled, but should only be triggered after the **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, it causes the old instance to be torn down before the entire system finishes changing routes using the new instance; this leads to traffic loss. PR1402382
- On devices running Junos OS, with scaled MPLS labels used, when the system is already running with high load, allocation of inefficient labels might cause even higher CPU utilization at 100% for hours. The issue might affect traffic. PR1405033
- When the **sr-mapping-client** statement is configured in IS-IS segment routing, the LDP route might not be present in inet.3 and routing-instance.inet.3, and also an invalid I/O label might be advertised in the LDP database. PR1416516
- The LDP transit egress route for a BGP route has an indirect next hop. In NSR and GRES scenario, after Routing Engine switchover, in some cases, LDP might fail to receive 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 might become unusable and the LDP transit egress route will get deleted. It could cause BGP sessions to go down and cause traffic drop. PR1420103
- LDP route metric might not match IGP route metric even with **ldp track-igp-metric** configured. PR1422645

Network Management and Monitoring

- Updating the address of the Juniper Networks Inc. in the SNMP MIB CONTACT-INFO entry - "{ snmpModules 1 }". PR1336291
- The snmpd process leaks memory in the SNMPv3 query path and crashes.

  The issue is caused by a memory leak when the request PDU is dropped by SNMP when the **snmp filter-duplicates** configuration 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 free 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
- 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

- This is a minor enhancement to add a UI to copy files from Junos VM to host Linux. PR1341550
- In a filter list (input-list/output-list) scenario, when the filters in the same filter list refer to the same nested filter, the FPC might crash continuously. The issue results in traffic loss during FPC crash and reboot. PR1357531
- With Junos OS Release 17.3R3 on MX Series, on moving from the baseline configuration to an EVPN scaled (4000 VLANs) configuration with multihoming, the newly elected designated forwarder might take up to 90 seconds to resume forwarding BUM traffic. The time required for convergence is proportional to the scale used, so a lower scale incurs a smaller dark window. Workaround for faster convergence with high scale: Distributing the configuration across several FPCs can potentially bring down the BUM traffic drop from 90 seconds to a significantly lower value. PR1362934
- There are multiple failures when events such as node reboots, ICL flaps, and ICCP flaps occur. Even with enhanced convergence configured, there is no guarantee that subsecond convergence will be achieved. PR1371493
- In a Layer 3 VPN network with large-scale prefixes, if the peer PE devices is other vendor's router configured, **per-prefix label**, all FPCs might restart after Layer 3 VPN routes churn multiple times. PR1398502
- In some cases, PS interfaces over RLT might be shown as up but they might not pass traffic. Log messages reporting an ASIC error and a chassis alarm reporting hard FPC errors may also be seen. PR1400269
- In some cases, the status bit of the RPF next hop appears as disabled when it should have been enabled. The trigger for the issue is not known yet. PR1404240
- Heap memory leak is seen during login / logout test-dhcpv4/v6 subscribers fails to login after multiple iterations. PR1442770

Routing Policy and Firewall Filters

- The rpd might crash during the policy configuration changes. PR1357802
- If a **policy option** with only the conditions **from route-distinguisher** and **then next-hop a.b.c.d** is applied to BGP, next hop for routes in the inet.0 might be set to the 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 Junos OS command **show bgp summary** does not show the BGP establised state. If the BGP state is not an ESTABLISHED 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 displayed properly. The issue happens for IPv4 BGP sessions only, on IPv6, all the BGP states are seen as default. PR600308
- In rare cases, the rpd might generate a core file with error **rt_notbest_sanity: Path selection failure on ...**" The core file is soft, which means there should be no impact to traffic or routing protocols. PR946415
- Junos OS shows an obsolete **Session description** in the output of **show route multicast extensive** for several multicast registry addresses. PR1022288
- **JTASK_SCHED_SLIP** for rpd might be seen on doing restart routing or ospf protocol disable with scaled BGP routes on an MX104 router. PR1203979
- LDP OSPF are 'in-sync' state and the reason observed for this is **IGP interface down** with LDP synchronization enabled for OSPF. **user@host> show ospf interface ae100.0 extensive Interface State Area DR ID BDR ID Nbrs ae100.0 PtToPt 0.0.0.0 0.0.0.0 0.0.0.01Type: P2P, Address: 10.0.60.93, Mask: 255.255.255.252, MTU: 9100, Cost: 1050Adj count: 1Hello: 10, Dead: 40, ReXmit: 2, Not StubAuth type: MD5, Active key ID: 1, Start time: 1970 Jan 1 00:00:00 UTCProtection type: NoneTopology default (ID 0) -> Cost: 1050LDP sync state: in sync, for: 00:04:03, reason: IGP interface downconfig holdtime: infinity** According to the current analysis, **IGP interface down** is observed as the reason because although LDP notified OSPF that LDP synchronization was achieved, OSPF was not able to take note of the LDP synchronization notification, because the OSPF neighbor was not up yet. PR1256434
- In an IS-IS and IPv6 scenario, rpd might crash when the neighbor router restarts and causes routes churn. PR1312325
- The rpd might crash and generate core files if distributed IGMP (Internet Group Management Protocol) is configured. PR1314679
- On a dual Routing Engine system with graceful Routing Engine switchover (GRES) and graceful restart enabled, if Bidirectional Forwarding Detection (BFD) with the **hold-down-interval** option is enabled on an external BGP peer, this BGP peer might stay in idle state after a Routing Engine switchover. PR1324475
- BGP I/O threading was 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
- When 32,000 SR-TE policies are configured at once, scheduler slips might occur during the configuration. PR1339829
- There are scenarios where an application allocates and caches next-hop templates. This causes next-hop template cache to grow continuously. But when the application clears their local cache, then memory is freed to next-hop template cache. But the next-hop template cache does not have the code to shrink the cache and free memory back. So the next-hop template memory is trapped in the cache and cannot

be used for other purposes. But if same BGP routes and nexthops come up again, they will reuse the templates from cache and not consume additional memory. PR1346984

- With BFD configured on an aggregated Ethernet (ae-) interface, if you disable and then enable the ae- interface, then that AE interface and the BFD session might not come up. PR1354409
- It is possible for a GNF with rosen6 multicast to display stuck KRT queue entries after recovery from a dual Routing Engine reboot at the BSYS. PR1367849
- When the loopback interface is configured in a logical system and Routing Engine-based micro BFD is configured to use the loopback address as the 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
- In Junos OS Release 18.4R1, RIB learning rate has degraded from anywhere between 10% and 18% on different platforms. For vale it seems to be 18%, whereas for all MX Series routers it is lesser than 10%. The root cause analysis is not completed and it is risky to include it in Junos OS Release 18.4R1. PR1383371
- At scale, a GNF with PS over RLT and multiple MPCs might show BFD flap at recovery. PR1386574
- On all devices running Junos OS, with GRES and nonstop routing (NSR) enabled, if Routing Engine switchover is executed, the Border Gateway Protocol (BGP) peers in the new master Routing Engine might flap due to hold-timer expiry after GRES. PR1390113
- In a BGP scenario with multipath enabled, if applying an import or export policy of IPv6 routes with a IPv4 next hop to a BGP neighbor, the rpd might crash continuously. PR1390428
- If an import policy is applied to a BGP neighbor and the policy has an indirect IPv4 next-hop for IPv4 and IPv6 routes (IPv6 routes resolved over IPv4), when a BGP unresolved route is withdrawn, rpd crash might be seen. PR1391568
- The **as-path-group** configuration is limited in scale. With 10,000 lines, scheduler slips are seen, impacting the other work rpd is doing such as protocol keep-alives. To avoid the scheduler slips (CPU exhaustion), change how the **as-path-group** is structured. The issue occurs due to two factors: the number of as-path statements under the **as-path-group** and the wild cards in each of these. In this PR, there is a new Junos command introduced: **set policy-options asregex-optimize**. The default feature is **no-optimize**. PR1396344
- It is possible that under certain scenarios when the legacy-redirect-ip-action the existing BGP routes advertised might not be refreshed. Because of this the routes might still contain communities not aligned with the configured **legacy-redirect-ip-action** option. Clear routes as described in workaround. PR1396787
- Users that replace simple VLAN interfaces with PS over RTL might notice an increase in FPC CPU usage. This is in keeping with the increased processing and resources needed to support these types of interfaces, which are similar in this regard to that of an AE interface. PR1396925
- When the multicast-only fast reroute (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 at the Packet Forwarding Engine level. PR1399457

- In a multicast routing scenario using PIM, if you are configuring a static route with **qualified-next-hop** for multicast source, rpd process might crash. This is because **qualified-next-hop** points to the GF_DLI (Gateway Family Data Links) address, which the PIM is unable to process, resulting in the crash. PR1408443
- In BGP with the indirect next-hop scenario, if unicast reverse-path forwarding (unicast RPF) is enabled, and then you enable BGP multipath, a background job loop might be formed and the CPU utilization of the rpd process might be stuck at 100%. PR1414021
- Change in route selection process. To select the better route between a non-BGP and BGP route, if you are at Step 7 of the route selection process (Understanding BGP Path Selection), then the BGP route is always the better one.. PR1415468
- In BGP multipath scenario with labeled-unicast (LU) enabled, if **no-propagate-ttl** is configured, the rpd might crash if BGP LU route's ttl action is changed after which it does not match BGP multipath cache. PR1425173
- In an MVPN scenario, the rpd might crash while removing multicast routes that do not have an associated (S,G) state or activating the **accept-remote-source** configuration on PIM upstream interface. PR1426921

Subscriber Access Management

- In a subscriber scenario, the authd might crash multiple times because of a memory corruption issue. PR1402012
- The authd addresses too quickly before jdhcpd can completely clean up the old subscriber which flooding error log . The log such as :**jdhcpd: %USER-3-DH_SVC_DUPLICATE_IPADDR_ERR: Failed to add 10.1.128.3 as it is already used by 1815** . PR1402653

User Interface and Configuration

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

VPNs

- The Multicast VPN MIB was not being properly compiled into the Juniper MIB package bundle. This might cause mib-jnx-mvpn.txt to be included as part of the Juniper Enterprise MIB set. PR1394946
- When the end-interface or backup-interface or 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

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: 18.4R2

Application Layer Gateways (ALGs)

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

Authentication and Access Control

- Push-to-JIMS now supports push auth entry to all online JIMS servers. PR1407371

Class of Service (CoS)

- The cosd process might crash while committing configuration through NETCONF. PR1403147
- Traffic drop occurs when deleting MPLS family or disabling an interface that has nondefault EXP rewrite rules. PR1408817

EVPN

- The EVPN implementation does not follow RFC-7432. PR1367766
- The rpd process crashes if the Autonomous-System (AS) is a deactivated 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 the EVPN peer goes down. PR1390965
- The rpd process might crash with EVPN type-3 route churn. PR1394803
- The BUM traffic might not be flooded in an EVPN-MPLS scenario. PR1397325
- IPv6 link-local address for the virtual-gateway address is marked as duplicate in EVPN. PR1397925
- When committing a configuration for adding 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 route withdrawal has 1 label. PR1399726
- ARP refresh functionality might fail in an EVPN scenario. PR1399873
- RPD core files upon Routing Engine switchover with scaled EVPN configuration. PR1401669
- The rpd crash because of the memory corruption in EVPN. PR1404351
- EVPN database and bridge MAC table are out of sync due to flapping of the interface. PR1404857
- The rpd might crash on a leaf node when handling the withdrawal of remote or local MAC addresses in an EVPN-VXLAN scenario. PR1405681
- The next hop is not cleaned up properly when one of the multihomed CE-PE links goes down. PR1412051
- Local l2ald proxy MAC+IP advertisements accidentally delete MAC+IP EVPN database state from remotely learned type 2 routes. PR1415277
- EVPN-MPLS single active :[EVPN/7] /32 host route always appears on non-DF PE if CNH is ON, **remote-ip-host-routes** has no effect. PR1419466
- rpd crash on backup Routing Engine after enable nonstop-routing with EVPN. PR1425687
- The device might proxy the ARP probe packets in an EVPN environment. PR1427109
- IP address is missing in **mac-ip-table** of the EVPN database but is present in the EVPN database when the CE interface has two primary IP addresses. PR1428581
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs. PR1429821
- Configuration is prevented from being applied on MX in subscriber scenario. PR1430360
- Incorrect MAC count with **show evpn/bridge statistics**.PR1432293

- Stale MAC addresses are present in the bridge MAC table in an EVPN-MPLS scenario. PR1432702
- Configuring ESI on a single-homed 25G port might not work. PR1438227

Forwarding and Sampling

- In an EVPN A-A scenario with an MX Series router or an EX Series switch acting as a PE device, flood next hops to handle BUM traffic might not get created or miss certain branches when the configuration is performed in a particular sequence. PR1377749
- The LSI binding for the IPv6 neighbor is missing. PR1388454
- Junos OS: Firewall filter terms named **internal-1** and **internal-2** being ignored (CVE-2019-0036). PR1394922
- In Junos OS Release 13.3R9.13, the firewall filter action "decapsulate gre", de-encapsulates GRE, IP-over-IP, and IPv6-over-IP, but in Junos OS Release 17.3R3.9, it only de-encapsulates GRE. PR1398888

General Routing

- Error drops in XM/MQSS fabric streams (q-node stats) are not accounted for in class-of-service fabric statistics. PR1338647
- Large-scale users' login and logout might cause mgd memory leak. PR1352504
- Traffic loss might be seen on the new master after the interface flaps followed by Routing Engine switchover in a VRRP scenario. PR1353583
- Packets might be dropped when they go through MX104 built-in interface. PR1356657
- MPC5E, MPC2E-NG, or 3E-NG might crash and restart during unified ISSU. PR1369635
- The dot1xd might crash when it receives an incorrect reply length from the authd. PR1372421
- Core files are seen in **ifinfo at pif_af_fe_info pif_af_ifd** when displaying af interface information. PR1373436
- MS-MPC might have performance degradation under scaled fragmented packets. PR1376060
- **NFX3/ACX5448:LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** is displayed during commit on configuration prompt. PR1376665
- MQSS errors might cause FPC restart. PR1380183
- The routes learned over an interface will be marked as "dead" next hop after changing the prefix length of an IPv6 address on that interface. PR1380600
- Traffic silently dropped because of an offline FPC in an MC-LAG scenario. PR1381446
- The unicast traffic from IRB interface toward LSI might be dropped due to Packet Forwarding Engine mismatching at egress processing. PR1381580
- PDT: MSE high CPU utilization for chassisd on BSYS, 20% st steady state. PR1383335

- The Virtual Chassis could not come up after upgrading to QFX5E platforms (TVP-based platforms for QFX5100 or QFX5200 switches). PR1383876
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent Major Alarm. PR1384435
- Subscriber connection setup is 30% lower than expected. PR1384722
- The rpd might crash when switchover is performed along with configuration changes being committed. PR1385005
- The device with more than five IP addresses configured in the DHCP server-group goes into Amnesiac mode after reboot. PR1385902
- The rpd end up with stuck krt queue might be seen in a VRF scenario. PR1386475
- Behavior of the CLI **set interfaces ams0 service-options session-limit rate <integer value>** command has changed. PR1386956
- Migrate from syslog API to Errmsg API - VM host messages on Junos OS. PR1387099
- On MX2000 platforms, backup CB's chassis environment status keeps 'Testing' after backup CB becomes online by removal/insert operation. PR1387130
- Chassisd process might have random memory corruption and will result in chassisd restart. PR1387338
- Some SFBs might go down when one of the PSMs in the chassis generates a bad output voltage that is out-of-range. PR1387737
- IPsec IKE keys are not cleared when delete/clear notification is received. PR1388290
- BBE SMGD generates core files if MTU is changed while subscribers are logged in on the physical interface. PR1389611
- The jnxFruState might show incorrect PIC state after replacing an MPC is replaced with another MPC with fewer PICs. PR1390016
- Traffic destined to VRRP VIP gets dropped as filter is not updated to the related logical interface. PR1390367
- Delete chassis redundancy will not give commit warning. PR1390575
- 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 a packet arrives from subscriber. PR1391932
- Third-generation FPC reboot loop because of having internal interface issues. PR1393643
- Junos OS enhancement configuration statement to modify mcontrol watchdog timeout. PR1393716
- IPv6 next-hop programming issue might be observed on QFX10000, PTX1000, and PTX10000 devices. PR1393937

- The FPCs might not come up during unified ISSU on MX10003. PR1393940
- CI-PR:Expected entries **UI_COMMIT_PROGRESS** are not getting populated while checking with Junoscript session for obtaining syslog output. PR1394780
- The l2ald process might crash during **commit check** for some specific configurations. PR1395368
- The minor alarm of "Bottom Fan Tray Pred Fail" might be incorrectly raised when the fan is at high speed on MX960. PR1395539
- Layer 3 gateway did not update ARP entries if IP or MAC quickly move from one router to another router in EVPN-VXLAN environment. PR1395685
- MPC7, MPC8, and MPC9 might not boot in MX Series Virtual Chassis. PR1396268
- The subscriber bindings might not be successful on QFX Series or EX Series platforms. PR1396470
- Adding IRB to bridge-domain with PS interface causes kernel crash. PR1396772
- The MS-MPC might generates core files when mspmand receives a non-syn packet of TCP. PR1396785
- Subscriber flapping may cause SMID resident memory leak. PR1396886
- Seeing **VMHost RE 0 Secure BIOS Version Mismatch** and **VMHost RE 1 Secure Boot Disabled** alarms. PR1397030
- mspmand core file is seen when committing configuration NAT pool changes to active NAT pool. PR1397294
- smid process memory leak and not coming down from 100%. PR1397643
- PFT MX10008: Inline-services Enabling the Flex-Flow-Sizing takes more than 12 minutes to move to steady state. PR1397767
- [jinsight] [generic_jinsight] show system errors active is not showing the error for MPC3E NG HQoS. PR1398084
- MPLSoUDP/MPLSoGRE tunnel might not come up on interface route. PR1398362
- High jsd or na-grpcd CPU usage might be seen even JET or JTI is not used. PR1398398
- IPsec tunnel cannot be established because the tunnel SA and rule are not installed in the PIC. PR1398849
- Incorrect timestamp is displayed in the jvision collector log file. PR1399829
- JET/PRPD incompatibility for the rib_service.proto field RouteGateway.weight from Junos OS Release 18.4R1 to 18.4R2 onward. PR1400563
- The mgd-api crashes due to memory leak. PR1400597
- Only one Packet Forwarding Engine could be disabled on an FPC with multiple Packet Forwarding Engines in error/wedge condition. PR1400716
- Config option forwarding-options enhanced-hash-key family mpls ether-pseudowire zero-control-word Does Not Take Affect in Junos Node Slicing. PR1400881

- The framed route beyond the first might not be installed in a DHCP subscriber management environment. PR1401148
- The authd might crash while restarting when you issue **show network-access requests pending**. PR1401249
- The command **show | compare** output on global group changes loses the difference context after a rollback or load update is performed. PR1401505
- The subscriber route installation failed because of improper installation of some interfaces states. PR1401506
- The TCP connection between ppmd and ppman might be dropped due to a kernel issue. PR1401507
- FPC core files are seen due to a corner case scenario (race condition between RPF, IP flow). PR1401808
- The **na-grpcd** log file is not rotated and keeps growing until Routing Engine is out of disk space. PR1401817
- JET authentication does not work for usernames and passwords of certain lengths. PR1401854
- Traffic loss is seen in IGMP subscribers after GRES. PR1402342
- The MPC might crash due to CPU overuse by dfw thread. PR1402345
- The device is in Amnesiac mode after ISSU with "mgd: error: configuration check-out failed" generate. PR1432664
- Some error logs might be seen on FPC when reading attempt from uninitialized memory location. PR1402484
- FPC might crash after MIC-3D-16CHE1-T1-CE-H is taken offline and brought back online. PR1402563
- DHCP subscriber cannot reconnect over dynamic VLAN demux interfaces due to RPF check failure. PR1402674
- Host outbound traffic might be dropped on MPC7, MPC8, and MPC9. PR1402834
- uncolored SRTE stats : MX: Observed rpd core files when a few colored LSPs were changed to uncolored LSPs. The core files are at **<<< #0 tag_cmp_tag (tag1=0x0, tag_label1=0x0, tag2=0x98b6628, tag_label2=0x98b6644) at ../../../../../../../../../src/junos/usr.sbin/rpd/lib/mpls/label_mgr/core/mpls_label.c:473 473 if (tag1->tagt_mtu != tag2->tagt_mtu) >>>**PR1403208
- Reported log variance might be incorrect if the PTP profile is changed from G.8275.2 to SMPTE or other multicast IP profile. PR1403219
- The smg service could become unresponsive when doing some GRE-related CLI operations. PR1403480
- The time synchronization through PTPoE might not work when Enhanced Subscriber Management is enabled on MX Series routers.PR1404002
- Continuous kernel crashes might be observed in backup Routing Engines or VC-bm. PR1404038

- With MS-MPC and MS-MIC service cards syslog messages for port block interim may show 0.0.0.0 as the private IP address and PBA release messages may show the NAT'd IP as the private IP. PR1404089
- The FPC might crash in a CoS scenario. PR1404325
- the repd continues to generate core files on Virtual Chassis-Bm when there are too many IPv6 addresses on one session (hit PR1384889). PR1404358
- Incorrect output of the assigned prefixes to the subscriber in the output of **show interface < dynamic demux interface>**PR1404369
- Configuring load override or load replace resets ANCP neighbors. PR1405318
- Voltage read failed for rail LTC3887-EA1-VDD0V9R2-CH0. PR1405787
- When using aggregated Ethernet bundle with active subscribers, FPC might crash if existing leg is replayed (after FPC restart). PR1405876
- NAT64 translation issues of ICMPv6 Packet Too Big message with MS-MPC/MS-PIC. PR1405882
- The FPC crash might be observed in an MS-MPC HA environment. PR1405917
- Fabric performance drop on MPC7/8/9E and SFB2-based MX2000 platforms. PR1406030
- The rpd might crash due to 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
- MX10003 cosmetic message: **ALARMD_CONNECTION_FAILURE: after 60 attempts craftd connect returned error: Connection refused**. PR1406952
- FPC might crash during the subscriber-related stresstests. PR1407285
- L2 VPN might flap repeatedly after the link between the PE and CE devices starts coming up. PR1407345
- The rpd might crash when a commit check is executed on LDP trace options. PR1407367
- Ephemeral DB might get stuck during commit. PR1407924
- Traffic forwarding failed when crossing VCF members. PR1408058
- openconfig-network-instance: network-instances support for IS-IS must be hidden unless supported. PR1408151
- 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 detected** after **request reboot vmhost routing-engine both** . PR1408480
- The MPCs 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-Service templates are not cleaned up. PR1409398

- MX-MPC2-3D-EQ and MPC-3D-16XGE-SFPP will now show "Exhaust A" temperature, rather than Intake temperature. PR1409406
- Telemetry: interface-set metadata needs to include the CoS TCP names in order to aid collector reconciliation with queue-stats data. PR1409625
- The CPU might be overused by jsd process in JET scenario. PR1409639
- The nonexistent subscribers might appear in the **show system resource-monitor subscribers-limit chassis extensive** output. PR1409767
- FPC might crash during next-hop change when using MPLS inline J-flow. PR1409807
- When using SFP+, the Interface optic output might be non-zero even after the interface has been disabled. PR1410465
- Traffic loss may be seen on MPC8E/MPC9E after you request one of the SFB2s that has gone offline to be brought back online. PR1410813
- Kernel replication failure might be seen if an IPv6 route next hop points to an **ether-over-atm-llc** ATM interface. PR1411376
- Packet Forwarding Engines heap memory leak might happen by frequent flapping of thousands of PPPoE subscribers. PR1411389
- Virtual route reflector may report DAEMON-3-JTASK_SCHED_SLIP_KEVENT error on some hypervisor or host machine because of NTP synchronization. Routing protocol may be impacted. PR1411679
- **file copy /var/tmp/file.name ftp://anonymous@< ip>/pub/** could not work properly after upgrade. PR1412033
- MX10003: The rpd crashes when the **switchover-on-routing-crash** does not trigger Routing Engine switchover and the rpd on the master Routing Engine goes into STOP state. PR1412322
- Junos PCC may reject PCUpdate/PCCreate message if there is a 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 due to an incorrect processing of a transient event. PR1413224
- During unified ISSU from Junos OS Release 16.1R4-S11.1 to Release 18.2R2-S1.2, CoS GENCFG write failures 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 ATM interfaces is broken in the Junos OS Release 17.2R1 onward. PR1413663
- The services load balance might not be effective for AMS if the hash key under the **forwarding-options** hierarchy is configured. PR1414109
- FPC crash might be observed if it reaches heap utilization limit. PR1414145
- NPC might not apply configured resource-monitor thresholds after NPC restart. PR1414650

- Firewall filters are not getting programmed into Packet Forwarding Engine. PR1414706
- The user might not enter configure mode as 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 when you run **show system subscriber-management route route-type <> routing-instance <>**. PR1415922
- Some IPsec tunnels might fail to pass traffic after GRES on an MX Series platform. 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 an IPv6 packet properly under NAT64 scenario. PR1417215
- Some subscribers might be offline when doing GRES or daemon restart. PR1417574
- Observed zero tunnel stats on the soft-gre tunnel. PR1417666
- The BGP session might flap after Routing Engine switchover. PR1417966
- 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** MX.PR1418461
- **sp-cleanup-timer** is not being honored when **lsp-cleanup-timer** is configured to be greater than 2147483647. PR1418937
- The reserved PPPoE ssession ID 65535 might also be assigned, which is in conflict with RFC 2516. PR1418960
- RX alarms are not set as according to the threshold value configured for the DCO Tunable optics. PR1419204
- A PPP session under negotiation might be terminated if another PPPoE client bears the same session ID. PR1419500
- CPU usage on Service PIC may spike while forming an IPsec tunnel in a DEP/NAT-T scenario. PR1419541
- A new tunnel could 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
- **bbe-mibd** memory leak causing daemon crash when having live subscribers and SNMP OIDs query. PR1419756
- In the scenario where the MX Series devices and the peer device both try to bring an IPsec tunnel up, so both sides are acting as an initiator, if the peer side does not answer the MX ISAKMP requests, the MX Series device can bring the peer-initiated tunnel down. PR1420293

- MX: PTP phase aligned but TE/cTE not good. PR1420809
- Failed to reload keyadmin database for **/var/etc/keyadmin.conf**. PR1421539
- **bbemg_smgd_lock_cli_instance_db** should not be logged as error messages. PR1421589
- MX Series Virtual Chassis: VCP port reports MTU value 9152 in the ICMP MTU exceeded message while the VCP port MTU is set to 9148. PR1421629
- The ps access interface is not marked ccc down on standby/non-designated PE. PR1421648
- **RPT_REG_SERVICES: RPM** syslogs are not getting generated after deactivating the aggregate interface. PR1421934
- 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
- While committing a huge configuration, the user might see the error **error: mustd trace init failed**. PR1423229
- **set forwarding-options enhanced-hash-key symmetric** is not effective on MX10003. PR1423288
- IP packet drop might be seen under Layer2 circuit scenario. PR1423628
- Traffic is dropped after FPC reboot with aggregated Ethernet member links deactivated by remote device. PR1423707
- 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
- The bbe-smgd process might crash after the command **"show system subscriber-management route prefix** is executed. PR1424054
- The port configured for 1-Gbps speed flaps after Routing Engine switchover. PR1424120
- The interface configured with 1-Gbps speed on JNP10K-LC2101 cannot come up. PR1424125
- [vMX]Continous disk error logs on VCP Console (Requesting switchover due to disk failure on ada1). PR1424771
- Interface with FEC disabled is flapping after Routing Engine mastership switchover. PR1425211
- In WAG scenario, soft-gre tunnel route lost after reboot/GRES or upgrade. PR1425237
- RPT_BBE_Regressions : Getting Unisphere-UpStream-Calc-Rate as 0 while verifying L2BSA RADIUS accounting stop packets after performing GRES. PR1425512
- All interfaces creation failed after NSSU. PR1425716
- IFL Targeting: 18000 phantom distributed interfaces are displayed for aggregated Ethernet interface with the targeted distribution enabled on it, when there are no active subscribers. PR1426157
- Interfaces might come to down after device reboots. PR1426349

- PEMs lose DC output power load sharing after PEM power-off and power-on operation on MX Series. PR1426350
- Traffic loss might be seen when multiple IPsec tunnels are established with the remote peer. PR1426975
- Traffic might not flow through MACsec interface even after an unsupported cipher-suite is removed. PR1427294
- When broadband edge PPPoE and DHCP subscribers coming up over Junos fusion satellite ports are active, **commit full** and **commit synchornization full** commands fail. PR1427647
- When installing YANG package without the **proxy-xml** configuration, the CLI environment did not work well. PR1427726
- The subscriber IP route may get suck in bbe-smgd if the subscriber IP address is the same as the local IP address. PR1428428
- PTSP subscriber stuck in configured state. Auto-clear-timer did not work as well. PR1428688
- Incorrect IGMP statistics for dynamic PPP interfaces. PR1428822
- L2TP subscriber and MPLS Pseudowire Subscriber volume accounting stats value remains unchanged after ISSU. PR1429692
- Destination unreachable counter was counting up without receiving traffic. PR1431384
- During the stresstests, bbe-smgd process might crash on backup Routing Engine when performing GRES. PR1431455
- The bbe-smgd might crash if subscribers are trying to log in or log out and a configuration commit activity happens at the same time. PR1431459
- Allow installation of three identical framed-routes in the same routing-instance. PR1431891
- MX10003 - **PEM not present** alarm raised when minimum required PEM exist in the system. PR1431926
- RSI & RSI brief should not include **show route forwarding-table** when Tomcat enabled. PR1433440
- Collected service statistics all 0 after ISSU for MPC2. PR1433589
- Lawful intercept for subscriber traffic is not programmed in Packet Forwarding Engine if it is activated by Access-Accept. PR1433911
- 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
- MPC7, MPC8, MPC9, MX10003 MPC, EX9200-12QS, EX9200-40XS line card might crash in a scaling setup. PR1435744

Infrastructure

- SNMP OID IFOutDiscards are not updated when drops increase. PR1411303
- The traffic to the NLB server might not be forwarded if the NLB cluster works on multicast mode. PR1411549

Interfaces and Chassis

- Constant dcpfe process crash might be seen if you are 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 may lead all FPCs to reboot. PR1386768
- DCD core files are seen after FPC restart if channelized interfaces are configured. PR1387962
- All DPCs might crash while adding or deleting a logical interface from the aggregated Ethernet bundle. PR1389206
- Decoupling of L2 logical interface configuration from bridge domain or EVPN configuration. PR1390823
- The dcd memory leak might be seen when committing 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 sub-interface 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 cause interface flapping during commit. PR1402122
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces misleading error message. PR1402606
- The **targeted-broadcast** statement does not work on an IRB interface. PR1404442
- The subscriber may not access the device due to the conflicting 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 an invalid ETH-LM request for unsupported outgoing logical interface. PR1420514
- Invalid speed value on an interface might cause other interface configuration loss. PR1421857
- The syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** upon LFM related config commit on aggregated Ethernet interfaces. PR1423586
- The cfmd might crash on DPCE. PR1424912

- The logical interfaces in EVPN routing instances might flap after committing configurations. PR1425339
- **flexible-queuing-mode** is not working on MPC5E of VC member1. PR1425414
- Upgrade from releases before Junos OS Release 17.4R1 to releases having PR-1425804 fix results in cleanup of existing ECFM PM-history and PM-sessions restarts freshly with MI index as 1. PR1425804
- CFM message flooding. PR1427868
- The vrrpd process might crash after deleting VRRP sessions for several times. PR1429906

Layer 2 Features

- The rpd crashes after an 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 toward all local CE-facing interfaces. PR1406807
- Broadcast traffics might be discarded in a VPLS local-switching scenario. PR1416228
- Commit error is seen but the commit is processed if adding more than o. PR1420082

Layer 2 Ethernet Services

- The SNMP query on LACP interface might lead to lacpd crash. PR1391545
- On EVPN setups, incorrect destination MAC addresses starting with 45 might show up when the **show arp hostname** command is used. PR1392575
- Log messages **dot1xd[]: task_connect: task ESP CLIENT:...: Connection refused** might be reported in Junos OS Release 17.4 or later. PR1407775
- Packets might be dropped if the traffic is forwarded on an LT interface. PR1410970
- The IRB interface might flap after configuration change is committed on any interface. PR1415284
- The IPv6 neighbor might become unreachable after the primary link goes down in a VPLS scenario. PR1417209
- jdhcpd becomes aware about some of the existing configurations only after 'commit full' or jdhcpd restart. PR1419437
- Change the nd6 next hops to reject next hop once L2 interfaces gets disassociated with IPv6 entries. PR1419809
- The jdhcpd process might consistently run at 100% CPU and not provide service if the **delay-offer** is configured for the DHCP local server. PR1419816
- jdhcpd daemon might crash during continuous stress test. PR1421569

MPLS

- DSCP bit marking of LSP self-ping is not compliant with rfc7746. PR1371486
- The rpd might crash on backup Routing Engine after switchover. PR1382249

- A 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 an LDP route with an indirect next hop is deleted. PR1398876
- The Layer 2 circuit 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 in BGP-LU with egress protection while committing configuration changes. PR1412829
- 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
- Rpd memory might leak when RSVP LSP is cleared/re-signaled. PR1415774
- LDP routes might flap if committing any configuration changes. PR1416032
- Traffic might be silently discarded due to a long LSP switchover duration in an RSVP-signaled LSP scenario. PR1416487
- Bad length for Sub-TLV 34 (RFC 8287) in MPLS Echo Request. PR1422093
- Bypass dynamic RSVP LSP tears down too soon when being used for protecting LDP LSP with the **dynamic-rsvp-lsp**. PR1425824
- mpls ping sweep stops working and the CLI stops responding. PR1426016
- MPLS LSP auto-bandwidth statistics miscalculations may lead to high bandwidth reservation. PR1427414
- When MBB for P2MP LSP fails, it is stuck in the old path. PR1429114
- MPLS ingress LSPs for LDP link protection are not coming up after of MPLS is disabled/enabled. PR1432138

Network Management and Monitoring

- The sub-agent such as mib2d might crash and restart after the AGENTX session timeout between master(snmpd) and sub-agent. PR1396967
- Child link missed from mib id dot3adAggPortAttachedAggID (OID - 1.2.840.10006.300.43.1.2.1.1.13). PR1410439
- The snmp query might not get data in scaled L2circuits environment.PR1413352
- Syslog match filtering does not work if a single line of **/etc/syslog.conf** is more than 2048 bytes. PR1418705

Platform and Infrastructure

- The kernel and ksyncd generate core files after dual CB flap at **rt_nhfind_params: rt_nhfind() found an nh different from that onmaster 30326**. PR1372875
- Jlock hog might be reported at restart routing. PR1389809
- Individual command authorization might cause mgd crash. PR1389944
- Traffic is dropped when passing through MS-DPC to MPC. PR1390541
- MX: RFC2544 is not functioning as expected due to platform validation getting skipped for the MX Series device (chassis based boxes). PR1396751
- RVT interface might flap. PR1399102
- In a scaled scenario (500 TWAMP control sessions and 500 TWAMP test sessions), a few TWAMP connections might fail to establish. PR1399547
- Syslog error messages: **[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_pfe:3211): pfe_id 0 cchip 0[LOG: Err] COS_HALP(cos_halp_get_fabric_stats_per_pfe:3272): No PFE found for pfe_id_start 0** . PR1402377
- MAP-E some ICMP Types cannot be encapsulated or de-encapsulated on SI interface. PR1404239
- Some files are missing during 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 bridgedomain and IP/MPLS network. PR1407200
- 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/MX-FPC card and MPC card on egress PE router in L3VPN. PR1409523
- Junos OS: Insufficient validation of environment variables in telnet client might lead to stack-based buffer overflow (CVE-2019-0053). PR1409847
- 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 a MIC is pulled out when this MIC is booting up. PR1414816
- op url command cannot run a script with libs from **/config/scripts**. PR1420976
- ARP request is not replied to although **proxy-arp** is configured. PR1422148
- show jnh trap-info with incorrect LU instance caused a crash and generated core files on FPC. PR1423508
- The native VLAN ID of packets might fail to be removed when leaving out. PR1424174
- The policer bandwidth might be incorrect for the aggregate interface after activating the configuration statement **shared-bandwidth-policer**. PR1427936

- Pre-fragmented ICMP IPv4 packets might fail to arrive at the destination. PR1432506
- Enabling 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
- A certain combination of allow and deny commands does not work properly after Junos OS Release 18.4R1. PR1438269

Routing Policy and Firewall Filters

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

Routing Protocols

- BGP might not advertise routes on the existing BGP peer after a Layer 3 VPN instance is added. PR1237006
- The VRF static route might not be exported when **route-distinguisher-id** is used on RR in a BGP Layer 3 VPN scenario. PR1341720
- 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 the backup Routing Engine if **proxy-macip-advertisement** is configured on an 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
- In rare cases rpd might crash after Routing Engine switchover when BGP multipath and Layer 3 VPN **vrf-table-label** are configured. PR1389337
- BGP IPv6 routes with IPv4 next hop causes rpd crash. PR1389557
- The ppmd on the Routing Engine might run with high CPU utilization after Routing Engine switchover. PR1392704
- Rpd core files on the backup Routing Engine during neighborship flap when using **authentication-key** with size larger than 20 characters. PR1394082
- Snoop-pseudowires enabled MCSNOOPD at an H-VPLS hub PE might drop an LSI for the spoke neighbour pseudowire off the control NH for IGMP query flooding upon this pseudowire active->standby->active transition followed by mcsnoopd restart at the hub. PR1394213
- The best and the second-best routes might have the same weight value if BGP PIC is enabled. PR1395098
- BGP DMZ LINK BANDWIDTH - not able to aggregate bandwidth, when applying the policy. PR1398000

- The rpd soft core files and inappropriate route selection might be seen when Layer 2 VPN is used. PR1398685
- The rpd process might crash in a BGP setup with NSR enabled. PR1398700
- Junos OS: BGP packets can trigger rpd crash when BGP tracing is enabled. (CVE-2019-0019) PR1399141
- The UHP behavior is not supported for LDP to SR stitching scenario. PR1401214
- There might be unexpected packet drops in MoFRR scenario if the active RPF path is disabled. PR1401802
- The rpd might crash when BGP **add-path send** is configured and NSR is enabled. PR1401948
- The rpd might be stuck at 100% 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
- Sometimes when a new logical router is configured, logical router core files might be seen on the system. PR1403087
- The rpd memory leak might be seen in IS-IS segment routing scenario. PR1404134
- Extended traffic loss might be seen after link recovery when source packet routing is used on OSPF P2P links. PR1406440
- IGMP join through PPPoE sub not propagated to upstream PIM. PR1407202
- M Series, MX Series, QFX Series: mcsnoopd core files generated immediately after the commit change related to EVPN-VXLAN configuration. PR1408812
- SID label operation might be performed incorrectly in an 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 the **vrf-import/vrf-export** configuration. PR1413686
- Dynamic routing protocol flapping with VM host Routing Engine switchover on NG-RE. PR1415077
- The IS-IS-SR route sent by the mapping server might be broken for ECMP. PR1415599
- Route information might be inconsistent between the RIB and OSPF databases when using the OSPF LFA feature. PR1416720
- Junos OS: OpenSSL Security Advisory [26 Feb 2019]. PR1419533
- A memory leak in rpd might be seen if source packet routing is enabled for the IS-IS protocol. PR1419800
- IPv6 IS-IS routes might be deleted and not be reinstalled when the MTU is changed at the logical interface level for family inet6. PR1420776
- The rpd might crash in a PIM scenario with **auto-rp** enabled. PR1426711
- The rpd might crash while handling the withdrawal of an imported VRF route. PR1427147
- The rpd might generate core files due to improper handling of graceful restart stale routes. PR1427987
- RPD might crash with OSPF overload configuration. PR1429765

Services Applications

- ms- used for IPSEC PIC is listed in show services ha detail as standby, cosmetic issue. PR1383898
- The spd might crash when **any-ip** is configured in the **from** clause of the NAT rule with the static translation type. PR1391928
- **SPD_CONN_OPEN_FAILURE: spd_svc_set_summary_query: unable to open connection to si-0/0/0 (No route to host)** PR1397259
- IP ToS bits are not copied to the outer IPsec header. PR1398242
- Invalid Layer 4 checksum might be observed in 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 to the access line information between ICRQ and PPPoE messages. PR1404259
- The stale si- logical interface might be seen when L2TP subscribers with duplicated prefixes or framed-route log in. PR1406179
- The kmd process might crash on MX Series and ACX Series platforms when IKEv2 is used. PR1408974
- [technology/subscriber_services/jl2tpd] [all] RPT BBE Regressions : ERA value does not match configured values while verify new ERA settings are reflected in messages log. PR1410783
- jpppd core files 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 a 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
- Traffic gets dropped when the end behind NAT is the responder. PR1435182

Software Installation and Upgrade

- JSU might be deactivated from FPC in case of power cycle. PR1429392

Subscriber Access Management

- 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
- Adding a firewall filter service using the **test aaa** command causes a crash in dfwd. PR1402051
- 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

- Authd telemetry: Linked pool head attribute is incorrect for single pools. PR1413293
- CoA-NACK is not sent when peforming negative COA Request tests by sending incorrect session ID. PR1418144
- Subscribers might not be able to re-login in Gx-plus provisioning scenario. PR1418579
- PPPoE session might be disconnected when LI attributes are received in access-accept with invalid data. PR1418601
- Address allocation issue with linked pools when using linked-pool-aggregation. PR1426244
- RADIUS authentication server might always be marked as DEAD. PR1429528

User Interface and Configuration

- The **show configuration** and **rollback compare** commands cause high CPU usage. PR1407848

VPNs

- 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 an MVPN+NSR scenario. PR1392792
- Downstream interface is not removed from multicast route after getting PIM prune. PR1398458
- Routes with multiple communities being rejected in inter-AS NG-MVPN scenario. PR1405182
- The multicast traffic drop might be seen when **static-umh** is configured in NGMVPN scenario. PR1414418
- The rpd might crash in rosen MVPN scenario when the 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: 18.4R1

Application Layer Gateways (ALGs)

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

Authentication and Access Control

- MAC move might occur in DHCP security scenario. PR1369785
- IPv4 or IPv6 DHCP-security client entries will be recorded on trusted ports as well. PR1390676

Class of Service (CoS)

- The 802.1P rewrite might not work on inner VLAN. PR1375189

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

EVPN

- EVPN/VXLAN: MAC entry is incorrectly programmed in the Packet Forwarding Engine, leading to some traffic being silently dropped or discarded. PR1231402
- MPLS label leak leads to label exhaustion and rpd process crash. PR1333944
- EVPN type-5 route might be lost if **chained-composite-next-hop** command is configured. PR1362222
- The l2ald memory might cross the threshold in an EVPN scenario. PR1368492
- Proxy ARP might not work as expected in an EVPN environment. PR1368911
- The rpd might crash in EVPN scenarios when configuring EVPN. PR1369705
- EVPN active or active multi homed PE device occasionally prefers to route to a directly connected prefix using LSPs toward the multi homed peer instead of going directly out the IRB interface (which is up). PR1376784
- The RA packets might be sent out without using the configured virtual gateway address. PR1384574

Flow-based and Packet-based Processing

- PIM register message might be dropped on SRX Series devices. PR1378295

Forwarding and Sampling

- Junos OS allows firewall filters with the same name under **[edit firewall]** and **[edit firewall family inet]** hierarchy levels. PR1344506
- L2ald crashes when trying to adjust mac-table-size configuration. PR1383665
- The filter counter is not written to the accounting file when accounting is enabled on the bridge firewall filter. PR1392550

General Routing

- TACACS access does not work after upgrade. PR1220671
- Routing Engine and Packet Forwarding Engine out-of-sync errors are seen in syslog. PR1232178
- The mspmand process might generate a core file in rare conditions due to 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 routers, the **show chassis led** command should not be displayed in possible completions of the **show chassis** command. PR1268848
- Flexible PIC concentrator (FPC) crash/reboot is observed when bringing up about 12,000 Layer 2 Bit Stream Access (L2BSA) subscribers simultaneously. PR1273353

- Error messages might be seen if flapping the aggregated Ethernet interface hosted on MPC-3D-16XGE card. PR1279607
- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/mobiled. PR1284625
- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/mspmand. PR1284643
- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/mspsmd. PR1284654
- PPPoE canot dial in due to all PADI dropped as "unknown iif" when the aggregated Ethernet configuration is deactivated or activated. PR1291515
- Wrong packet statistics are reported in 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 overriride conditions. PR1313009
- Migrate from syslog API to Errmsg API - /src/junos/usr.sbin/subinfo. PR1327262
- Migrate from syslog API to Errmsg API - /src/junos/usr.sbin/aaad. PR1327266
- Migrate from syslog API to Errmsg API - /src/junos/usr.sbin/smihelperd. PR1327271
- Tc_count counters in filter with the **scale-optimized** command are not incrementing. PR1334580
- With certificate hierarchy, where intermediate CA profiles are not present on the device, in some corner cases, the PKI daemon can become busy and stop responding. PR1336733
- AI-script does not get automatically upgraded unless it is manually done after a Junos OS upgrade. PR1337028
- Routing Engine does not have MAC map for MAC type 7. PR1345637
- Additional **show** commands are called when the **request support information** command is issued. PR1346129
- 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
- Migrate from syslog API to Errmsg API - /bbe-svcs/smd/plugins/cos/. PR1353179
- Some of the inline service interfaces cannot send out packets with the default bandwidth value (100 Gbps). PR1355168
- Chassis alarm is not reflecting the correct state when INP0 and INP1 have AC voltage out-of-range. PR1355803
- The mpls-ipv4 template does not have correct src AS and dst AS as 4294967295 src Mask and DstMask as 0 after adding the mpls-flow table size on the fly. PR1356118
- Link stays up unexpectedly on MX204 with copper cable removed. PR1356507

- MPC/FPC might be unable to reply request messages to the Routing Engine in a high subscriber scale scenario. PR1358405
- **show chassis ethernet-switch** on PTX10000. PR1358853
- The **show chassis fpc** command output might show "Bad Voltage" for FPC powered off by configuration or CLI command after the command **show chassis environment fpc** is executed. PR1358874
- Bbe-smgd restarts unexpectedly while performing graceful Routing Engine switchover (GRES). PR1359290
- PluginExit() function is never called. PR1359610
- FPC core file might be observed after GRES switchover. PR1361015
- IP over VPLS traffic is affected by EXP rewrite rule on the core-facing MPLS interface. PR1361429
- The MX Series router functioning as a BNG does not generate ESMC/SSM Quality Level failed snmp trap. PR1361430
- Rpd struck at 100 percent after clear bgp neighbor operation. PR1361550
- Migrate from syslog API to Errmsg API;usp/usr.sbin/nsd/common/nsd_tpm.c. PR1361986
- Spontaneous bbe-smgd core file might be seen on the backup Routing Engine. PR1362188
- The MS-MPC might reset continuously on MX Series platforms. PR1362271
- M/Mx: Traffic loss of 1 percent is seen during GRES phase of unified ISSU from 17.3-20180527.0 to17.3-20180527.0. PR1362324
- Executing **show route prefix proto ip detail** during route churm in a route scale scenario might lead to FPC crash. PR1362578
- The inline-J-Flow sampling configuration might cause FPC crash on MX Series platforms. PR1362887
- MX-VC: Request to record VCCP heartbeat state change in syslog by default. PR1363565
- xmlproxyd for internal interfaces is reporting uint32 instead of uint64. PR1363766
- The multicast route update might get stuck in KRT queue and the rpd might crash if rpd and kernel go out of sync. PR1363803
- FPM board is missing in SNMP MIB walk. 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/activating interfaces/filter/class-of-services configurations due to accessing stale memory entry. PR1364477
- Configuration commit might be delayed by 30 seconds. PR1364621
- AF's operational state moves to down state in a node virtualized environment where GNFs are connected through AF interface. PR1364921
- The traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with "Link-Layer-Down" flag set. PR1365263

- Default adapter type changed from E1000 to VMXNET3. PR1365337
- Traffic drops seen if training failure is seen on a line card for three of more planes. PR1365668
- MPC7E: ukern crash and FPC reboot with vty command **show agent sensors verbose**. PR1366249
- MS-MPC/MS-PIC might crash in NAT scenario. PR1366259
- MX150: Upgrade to Junos OS Release 18.1R1.9 fails. Installing package **nfx-2-routing-data-plane-1.0-0.x86_64** needs 76 MB on the file system. PR1366324
- Migrate from syslog API to Errmsg API - junos/lib/liboiu-ffp/. PR1366546
- The next hop of MPLS path might be stuck in hold state, which could cause traffic loss. PR1366562
- Snmp MIB walk for UDP flood gives different output statistics than CLI. PR1366768
- Syslog errors seen **LOG : Err] Failed to allocate 2 jnh-dwords for encap-ptr(ether-da)!,LOG: Err] gen_encap_common: jnh-alloc failed! 8**. PR1366811
- Offline of the fabric links of Packet Forwarding Engine 4 and Packet Forwarding Engine 5 is not supported. PR1367412
- The bbe-smgd process might crash during the authentication phase for L2BSA subscriber. PR1367472
- The **show system resource-monitor fpc** output might show a non existing Packet Forwarding Engine. PR1367534
- RTG interface status might be shown as incorrect status with **show interface**. PR1368006
- Multiple provisioning and deprovisioning cycles cause rdmd memory leak. PR1368275
- JSA10893: 2018-10 Security Bulletin: MX Series: In BBE configurations, receipt of a crafted IPv6 exception packet causes a denial of service (CVE-2018-0058). PR1368599
- RPD API **rt_nexthops_extract_gateway_convert_unnumbered_gf_dli()** rectification. PR1368855
- The **commit** or **commit check** might fail due to the error of **not having lsp-cleanup-timer without lsp-provisioning**. PR1368992
- 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 committing demux-related configuration. PR1370015
- The rpd might crash after Routing Engine switchover is performed or the rpd is restarted if interface-based dynamic GRE tunnel is configured. PR1370174
- Packet that exceed 8000 bytes might be dropped by MS-MPC in ALG scenario. PR1370582
- GMIC2 : SFP-1FE-FX optics does not come up on GMIC. PR1370962
- All the MX150 devices running VRRP on a LAN are stuck in master state. PR1371838
- BBE SMGD generates a core file on FPC restart. PR1371926

- FPC high CPU utilization or crashes occur during hot-banking condition. PR1372193
- SMGD generates a core file after essmd restart with reference to **mmf_ensure_mapped** (mmf=0xe8f0200, offset=4294967295, len=108) at **../src/junos/lib/libmmf/mmf.c:1972**. PR1372223
- Need a way to verify the session IDs above the 32-bit limit to check if this is working. PR1385237
- With very high scale l3vpn, traffic is dropped when egressing on an AF interface. PR1372310
- Image installation on SD fails with error **Unable to read reply from software add command to re1; error 1**. PR1372877
- The Routing Engine might crash after non-GRES switchover. PR1373079
- Core in ifinfo at pif_af_fe_info pif_af_ifd when displaying af interface information. PR1373436
- AOC Type Optics fail to initialize on MACsec TIC startup. PR1373572
- EDVT-GI-MIC2 : Interfaces do not come up for bidirection module SFP-100BASE-BX10-U and SFP-100BASE-BX10-D. PR1373795
- BOOTP packets might be dropped if BOOTP-support is not enabled at the global level. PR1373807
- LDP convergence delay might be seen after IGP metric change with **bgp-igp-both-ribs** command configured. PR1373855
- There is a vMX QoS performance issue in the Junos OS Release 18.3. PR1373999
- Cosmetic log **warning: [---] is protected, 'protocols ---' cannot be deleted** is seen after commit using **configure private** in a configuration with "protect" flag present. PR1374244
- 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
- Bbe-smgd generates a core file continously while deleting multicast group node from the tree. PR1374530
- PCE-initiated LSPs remain **Control status became local** after removing PCE configuration. PR1374596
- A few L2BSA subscriber logical interfaces are left behind in SMD infrastructure 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
- The bbe-smgd core file might be seen after doing GRES. PR1376045
- Interface optic output power is not zero when the port has been disabled. PR1376574
- CI: Not generating Power Supply failed trap. PR1376612
- Disabling OAM might cause the Broadband Edge daemon to crash. PR1377090
- Packets might be dropped on data plane in the inline J-Flow scenario. PR1377500
- MQTT keepalive timeout messages seen in case of slow JTI collectors. PR1378587

- After NAT64 router (with MS-MPC) translates an IPv6 fragment to IPv4 fragment, router is not inserting the right value in identification field of IPv4 header. PR1378818
- The ICMPv6 packets larger than 1024 might be dropped if **icmp-large-packet-check** is configured on IDS service. PR1378852
- Traffic might get silently dropped or discarded when CoS configuration is changed on a PS interface. PR1379530
- Protocol adjacency might flap and FPC might reboot if jlock hog happens. PR1379657
- Remove 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, and system resumes normally processing on its own. PR1380231
- CE_Customer: DT_BNG: ESSM model: rpd generates a core file during the fifth GRES, with reference to **task_kevent_udata_task (ev= <optimized out>) at ../../../../../../src/junos/lib/libjtask/base/platform/bsd/task_io_bsd.c:127**. PR1380298
- Encryption and decryption do not occur, because the Packet Forwarding Engine discards while testing that the group VPN member was established by using the authentication-method preshared key ASCII text. PR1381316
- Memory leak observed in MS-MPC card. PR1381469
- Subscribers not able to log in after double GRES, after reboot, or after configuration. PR1382050
- On Summit MX3ru for Junos OS Release 18.3R1 release ISSU fails if QSA is plugged in. PR1382126
- The MPC6E might crash while fetching PMC device states. PR1382182
- Flows are getting exported before the active timeout. PR1382531
- PFT MX10008 expected **inline-ipv4-export-packet-failures** is not listed in show services accounting error. PR1382873
- MAC addresses might disappear, if the interface MTU of EVPN PE device is changed. PR1382966
- The kmd crashes with a core file after bringing up IPsec connection. PR1384205
- CoS attachment might be mistakenly removed for DHCPv4 stack when DHCPv6 stack fails to be brought up for single-session dual-stack subscriber. PR1384289
- MBFD flaps because clksync congest the scheduler for 100ms. PR1384473
- CE_Customer: DT_BNG: Bbe-smgd generates multiple core files with reference to **bbe_mcast_vbf_dist_policy_service_encoder (params= <opyimized out>) at ../src/junos/usr.sbin/bbe-svcs/smd/plugins/mcast/bbe_mcast_policy_config.c:159**. PR1384491
- RPT_REG_SERVICES: The MPLS packets with more than eight labels will not be processed by J-Flow. PR1385790

- IPsec VPN traffic might fail when passing through MS-MPC of MX Series routers with CGNAT enabled. PR1386011
- Representation of memory units is changed from gigabytes (GB) to gibibytes (GiB) in the help string under the resource template hierarchy. PR1386516
- RBU_REGRESSIONS_SERVICES ::IPv4 and IPv6 VIP Routes are not withdrawn after aggregated Ethernet and VLAN with IRB flap. PR1386713
- RBU_Services_Regressions: SFLOW : Agent ID in **show sflow** command is displaying lo interface IP instead of fxp0 IP. PR1386890
- In case a LSP is locally configured without an explicit path ERO, the object remains empty in the PCRpt generated by PCC. PR1386935
- Uninitialized EDMEM[0x400094] Read (0x6db6db6d6db6db6d) logs are seen with sampling applied to a subscriber with routing-service applied. PR1386948
- When tracing is enabled, having a lot of trace-flags set could result in an rpd core file due to buffer overflow. PR1387050
- The pccd might crash when changing delegation-priority. PR1387419
- The bbe-smgd daemon 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
- FPC core file might be seen at **sensor_export_time_exceed_limit agent_health_monitor_data_reap** when Jinsight is configured. PR1388112
- Bbe-smgd does not respond to NS from SLAAC client on dynamic VLAN. PR1388595
- Incorrect values for flow packets/octets fields might be seen in inline J-Flow scenario. PR1389145
- The bbe-smgd process generates repeated core files and stops running as a result of long-term session database shared memory corruption. PR1388867
- IGMP group threshold exceed log message prints a wrong demux logical interface. PR1389457
- BFD flaps are seen on MX Series platforms with inline BFD. PR1389569
- MX204 - Excluding **speed** CLI option under the interface level. PR1389918
- Class of service adjustment-control-profile configuration for application DHCP tags does not get applied. PR1390101
- Delay in CLI output with second or more **show subscriber <> extensive** queries occur when the first session is sitting at -(more)- prompt displaying **show subscribers extensive**. PR1390762
- Trailing characters appear in the GNMI get API reply. PR1390967
- DT_BNG: DFW plug in NACKs DHCPv6/PPPoE requires ESSM subscriber re-login after ISSU. PR1391409

- The **routing-engine-power-off-button-disable** command does not work on MX204. PR1391548
- The bbe-smgd process might crash after committing configuration changes. PR1391562
- On MX Series routers serving as a DHCP server for dual-stack subscribers, BBE-SMGD process generates a core file. PR1391845
- On MX2000, fans start spinning at high speed upon inserting previously offlined FPC. PR1393256
- If FPGA on the new master CB has a specific hardware failure, the chassid might keep crashing after GRES switchover. PR1393884
- PFT MX10008: Inline-services enabling the Flex-Flow-Sizing takes more than 12 minutes to move to steady state. PR1397767
- The **show system errors active** is not showing the error for MPC3E NG HQoS. PR1398084
- Kernel core file occurs on vMX due to jlock assert. PR1398320
- High jsd or na-grpcd CPU usage might be seen even JET or JTI is not used. PR1398398
- The bbe-smgd process might generate a core file when executing **show pppoe lockout**. PR1398873
- FPC might crash after offline/online MIC-3D-16CHE1-T1-CE-H. PR1402563

High Availability (HA) and Resiliency

- Backup Routing Engine might go to db prompt after performing configuration remove and restore. PR1269383
- Observed **error: not enough space in /var on re1.** while doing unified ISSU upgrade from Junos OS Release 17.4-20180328.0 to Release 18.2-20180416.0. PR1354069
- VC-Bm cannot sync with VC-Mm when the Virtual Chassis splits the reforms. PR1361617

Interfaces and Chassis

- Aggregated Ethernet speed calculation changes according to 10 Gigabit Ethernet after post GRES. PR1326316
- Momentary dip in traffic occurs when a GRES is performed. PR1336455
- Native-vlan-id support on ps-interface. PR1352933
- The sonet interface will go down after enabling "keep-address-and-control" in L2VPN scenario. PR1354713
- The aggregated Ethernet interface might flap when the link speed of the aggregated Ethernet bundle is configured to oc192. PR1355270
- Approximately 50 percent of PPPoE subscribers (PTA and L2TP) and all ESSM subscribers are lost after ISSU during DT CST stress test. PR1360870
- Error messages like **ifname [ds-5/0/2:4:1] is chan ci candidate** are seen during a commit operation. PR1363536

- In case of MPLS , DMR packets are sent with different mpls exp bits if MX Series router receives CFM DMM packets with varying exp values on MPLS header. PR1365709
- In rare case, there might be L2TP subscribers stuck in terminated state. PR1368650
- The EOAM LTM messages might not get forwarded after system reboot in CFM scenario configured with CCC interface. PR1369085
- ISSU could be aborted at **Timed out Waiting for protocol backup chassis master switch to complete** with MX Series Virtual Chassis configuration. PR1371297
- The error **parse_remove_ifl_from_routing_inst() ERROR : No route inst on et-0/0/16.16386** is seen after restarting l2cpd daemon. PR1373927
- The dcd process might go down when **vlan-id none** is configured for the interface. PR1374933
- FTI logical interface VNI limits changed from (0..16777215) to (0..16777214). PR1376011
- Duplicate IP cannot be configured on both SONET (so-) interface and other interfaces. PR1377690
- Some error logs (Tx unknown LCP packet) might be reported by the bbe-smgd daemon on MX Series platforms. PR1378912
- Higher level OAM CFM between CE might not work in VPLS scenario. PR1380799
- The dcd restarted unexpectedly after committing a configuration with static demux interface stacking over ps interface. PR1382857
- The jpppd process might crash if the EPD value contains a format specifier. PR1384137
- DCD core can be seen after FPC restart if channelized interfaces are configured. PR1387962
- Interface-control thrashes and dcd does not restart after adding invalid demux interface to the configuration. PR1389461
- Decoupling of Layer 2 logical interface configuration from bridge-domain or EVPN configuration PR1390823

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 Routing Engine switchover. PR1205373
- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/lacpd. PR1284592
- ZTP infra scripts are not included for MX Series PPC routers. PR1349249
- Migrate from syslog API to Errmsg API:PPMD client LACP. PR1358599
- The DHCP leasequery message is replied to with an incorrect source address. PR1367485
- JSA10889 2018-10 Security Bulletin: Junos OS: The jdhcpd process crashes during processing of specially crafted DHCPv6 message (CVE-2018-0055). PR1368377

- The kernel core might happen by commit operation in rare condition. PR1369459
- 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

Layer 2 Features

- The traffic might not be transmitted correctly in a large-scale VPLS scenario. PR1371994

MPLS

- When minimum-bandwidth and bandwidth commands are present in the configuration, the bandwidth selection of the LSP is inconsistent. PR1142443
- JDI-RCT: Rpd core file is seen on master Routing Engine after performing restart chassisd. PR1352227
- 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 in BGP LU and LDP scenario. PR1366920
- RSVP authentication might fail between some Junos OS releases and causes 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 a Juniper Networks device to another device not supporting RFC 6424. PR1372924
- Rpd process eventually might crash after Routing Engine switchover with GRES/NSR enabled. PR1373313
- The traffic might not be load-balanced equally across LSPs with ldp-tunneling configured. PR1373575
- The rpd process might crash continuously if nsr-synchronization or all flag is used in RSVP traceoptions. PR1376354
- JSA10883: Junos OS: Receipt of a specifically crafted malicious MPLS packet leads to a Junos kernel crash (CVE-2018-0049). PR1380862
- Ingress LSPs go down due to CSPF failure. PR1385204
- Configured bandwidth 0 does not get applied on RSVP interface. PR1387277
- Bypass LSP is taking same SRLG colored path. PR1387497

Platform and Infrastructure

- MAC addresses are not learned on bridge-domains after XE/GE interface flap tests. PR1275544
- MQCHIP CPQ block should report major alarm. PR1276132
- Distributed multicast might not be forwarded to a subscriber interface. PR1277744
- **show igmp statistics** not including any statistics under interface aggregate for distributed multicast interfaces. PR1289415

- When chassis control restart is done with aggregated Ethernet and COS rewrite configuration, **Platform failed to bind rewrite** messages could be seen in syslog. PR1315437
- RLT subinterfaces are not reporting statistics. PR1346403
- lt- interface gets deleted with tunnel-services configuration still present. PR1350733
- Some linecards might crash in subscriber scenario enabled with distributed IGMP. PR1355334
- When **forwarding-class-accounting** command 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
- JSA10899 2018-10 Security Bulletin: Junos OS: Nexthop index allocation failed: private index space was exhausted through incoming ARP requests to management interface (CVE-2018-0063). PR1360039
- Select CLI functions are not triggering properly (set security ssh-known-hosts load-key-file, set system master-password). PR1363475
- Qmon sensors are not working with hypermode enabled. PR1365990
- Subscribers over aggregated Ethernet interface might have tail drops, which will affect the fragmented packets due to QXCHIP buffer getting filled up. PR1368414
- Forwarding is broken after adding protocol **evpn extended-vlan-id**. PR1368802
- The host outbound traffic might get dropped when the **class-of-service host-outbound-traffic ieee-802.1 rewrite-rules** command is configured. PR1371304
- Traffic might drop on new added interfaces on MX Series routers after unified ISSU. PR1371373
- The logical tunnel interface might be unable to send out control packets generated by Routing Engine. PR1372738
- JNH memory leaks in multicast scenario with MoFRR enabled. PR1373631
- Traffic traversing an IRB is not tagged with a VLAN if the packets go through an additional routing-instance. PR1377526
- FPC crash might be seen after FPC restarts. PR1380527
- lsi binding is missing upon nd6 entry refresh after l2ifl flap. PR1380590
- Packet drops on interface if the command **gigether-options loopback** is configured. PR1380746
- In certain Junos scenarios, DFWD memory corruption is seen due to large logical interface fstate messages. This can lead to log messages on dfwd traceoptions and occasionally DFWD core file. PR1380798
- Packet drops might be seen if the packet header is over 252 bytes. PR1385585
- RADIUS not working using management instance for IPv6 family. PR1391160
- The configuration through NETCONF session might fail. PR1383567
- L3VPN/ROSEN over PS over RLT: In Junos OS Release 18.4DCB after ifconfig goes down for PS logical interface, and its Link and Admin status are not going down as expected. PR1396335

Routing Policy and Firewall Filters

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

Routing Protocols

- Migrate from syslog API to Errmsg API;/src/junos/usr.sbin/ppmd. PR1284621
- Multihop eBGP peering session exchanging EVPN routes can result in rpd core file when BGP updates are sent. PR1304639
- The BGP session might be stuck with high BGP OutQ value after GRES on both sides. PR1323306
- The rpd might crash when BGP neighbor is flapping. PR1337304
- The bfd process memory leak might be observed if enabling multi-hop BFD session for a static route with multiple qualified-next-hop. PR1345041
- Rpd crash might be seen after executing Routing Engine switchover. PR1349167
- FPC might continuously crash on vMX platforms. PR1364624
- sBFD session flaps incrementally with 300 StaticSR clients configured with 100 ms as minimum-interval. PR1366124
- Static route gets unexpectedly refreshed on commit when configured with resolve configuration statement. PR1366940
- About 10 minutes of traffic loss is caused by BGP flap during MX Series unified ISSU. PR1368805
- TCP sessions might be taken down during Routing Engine switchover. PR1371045
- Route entry might be missing when IS-IS shortcut is enabled and MPLS link flaps. PR1372937
- SSH is not working if **[edit system services ssh hostkey-algorithms]** is set or in FIPS mode. PR1382485
- The rpd might crash after issuing operational command **show route detail** for RIP route. PR1386873
- Penultimate-hop router does not install BGP LU label, causing traffic to be silently dropped or discarded. PR1387746
- Next hop is not deleted by ukernel. However, the **delete** command is seen in rtsockmon. PR1389379
- The rpd process might crash when **rp-register-policy** is configured with more than 511 terms. PR1394259

Services Applications

- Selectively start ZLB Delay timer at the Packet Forwarding Engine for LAC tunnels. PR1338450
- L2TP Access Concentrator (LAC) tunnel connection request packets might be discarded on LNS device. PR1362542
- The L2TP subscribers might not be able to log in successfully due to the jl2tpd memory leak. PR1364774
- Accounting stop message is not sent to 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
- L2TP subscribers might be stuck in init state in a corner case. PR1391847

Subscriber Access Management

- The authd process might not be started after executing Routing Engine switchover on the backup Routing Engine without GRES enabled. PR1368067
- RADIUS VSAs, Actual-Data-Rate-Downstream, and Actual-Data-Rate-Upstream values are not compliant with RFC 4679. PR1379129
- CoA updates 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
- Log Message: authd: gx-plus: logout: wrong state for request session-id <xyz>. PR1384599
- Multiple IPv6 IANA addresses are assigned for one session in IPv6 PD binding failure scenarios. PR1384889
- Usage-Monitoring-Information AVP as part of PCRF gx-plus provisioning is causing service accounting activation. PR1391411

VPNs

- The rpd process might crash after configuration change in an L2VPN scenario. PR1351386
- EOAM group-down status does not work as expected. PR1361437
- In dual-homed next-generation MVPN, the receipt of type 5 withdrawal removes downstream join states for some routes. PR1368788
- In MVPN source site, a redundant environment primary site can generate type 5 routes for the sources from different sites without having real traffic, potentially causing an outage if the receiver PE devices accept those routes as preferable. PR1375716
- The rpd process crashes when LSP template for a provider tunnel is changed. PR1395353

SEE ALSO

Documentation Updates

IN THIS SECTION

This section lists the errata and changes in Junos OS Release 18.4R2 documentation for MX Series.

Subscriber Management Provisioning Guide

- The new topic, Subscriber Management RADIUS Dictionary Files, provides a link to the Juniper Networks RADIUS dictionary that is used by default with subscriber management for each supported release. The dictionary is updated only when software features that affect the file are added or changed. The dictionary is not updated for every Junos OS release.
- Starting in Junos OS Release 15.1, the *Broadband Subscriber Sessions User Guide* and the CLI Explorer incorrectly included information about the **show extensible-subscriber-services accounting** command. This command is not present in the CLI. Instead, you can use accounting profiles to collect statistics from the Packet Forwarding Engine for Extensible Subscriber Services Manager (ESSM) subscribers. See Flat-File Accounting Overview for information about accounting for ESSM subscribers.

Subscriber Management VLANs Interfaces Guide

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

SEE ALSO

New and Changed Features | **90**

Changes in Behavior and Syntax | **110**

Known Behavior | **119**

Known Issues | **123**

Resolved Issues | **140**

Migration, Upgrade, and Downgrade Instructions | **173**

Product Compatibility | **180**

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the procedure to upgrade Junos OS, and the upgrade and downgrade policies for Junos OS for the MX Series. Upgrading or downgrading Junos OS might take several minutes, depending on the size and configuration of the network.

Starting in Junos OS 18.3R1 release, FreeBSD 11.x is the underlying OS for all Junos OS platforms which were previously running on FreeBSD 10.x based Junos OS. FreeBSD 11.x does not introduce any new Junos OS related modifications or features but is the latest version of FreeBSD.

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

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

Basic Procedure for Upgrading to Release 18.4

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-x86-32-18.4R2.9-signed.tgz
  ```

- For 64-bit Routing Engine version:

  ```
  user@host> request system software add no-validate reboot
  source/junos-install-mx-x86-64-18.4R2.9-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-x86-32-18.4R2.x-limited.tgz
  ```

- For 64-bit Routing Engine version:

  `user@host>` **request system software add no-validate reboot *source*/junos-install-mx-x86-64-18.4R2.9-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 18.4 **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-18.4R2.9-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-18.4R2.9-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 18.4 **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 18.4

To downgrade from Release 18.4 to another supported release, follow the procedure for upgrading, but replace the 18.4 jinstall package with one that corresponds to the appropriate release.

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

For more information, see the Installation and Upgrade Guide.

SEE ALSO

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for NFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 for the NFX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

New and Changed Features

IN THIS SECTION

This section describes the new features or enhancements to existing features in Junos OS Release 18.4R2 for NFX Series devices.

New and Changed Features: 18.4R2

There are no new features or enhancements to existing features in Junos OS Release 18.4R2 for NFX Series devices.

New and Changed Features: 18.4R1

Virtual Network Functions (VNFs)

- **vSRX Support**—Starting in Junos OS Release 18.4R1, vSRX 3.0 is supported on NFX250 devices.

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 18.4R2 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.

SEE ALSO

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 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 NFX150 devices, the TCP and ICMP RPM probes take the best-effort queue of the outgoing interface, instead of the network control queue. As a workaround, configure a DSCP value such as nc1 for the RPM probes to take the network control queue. PR1329643

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
- On NFX250 NextGen devices running Junos OS Release 18.4R1, the memory values of vjunos0, flowd and ovs are as follows:

```
Component  S1E     LS1   S1           S2
Vjunos0     1.95 G  1.95 G 1.95 G 1.95 G
Flowd     2.02 G  2.02 G 2.02 G 2.02 G
OVS       4.10 G  4.10 G 4.10 G 4.10 G
```

  PR1366147
- Starting in Junos OS Release 18.4, NFX150 devices support two versions of disk layout. In the older version of the disk layout, you could upgrade or downgrade from Junos OS Release 18.4. With the new disk layout, a downgrade to releases later than Junos OS Release 18.4 is not possible. As a workaround, avoid operations that reformat the disk layout. PR1379983

SEE ALSO

New and Changed Features | **181**

Changes in Behavior and Syntax | **182**

Known Issues | **185**

Resolved Issues | **187**

Documentation Updates | **188**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.4R2 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
- 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 NFX250 devices, libvirt is hung due to which the console access to the device and JDM do not work. PR1341772
- Starting in Junos OS Release 18.3R1, the reboot time has increased for fpc0 and fpc1 interfaces on NFX150 devices.PR1355527

- On NFX250 NextGen devices running Junos OS release 18.4, the **show vmhost network nfv-back-plane** command output shows the **Link State/Admin State** as down. PR1375908
- On NFX150 devices, when the interface configuration has the encapsulation **flexible-ethernet-services** enabled on a 10G interface, traffic gets dropped. PR1425927

Routing Protocols

- When there is a static route and an OSPF route is active in the routing table for a specific destination network, a ping initiated to that destination network from the NFX device will fail. PR1438443

High Availability

- On an NFX150 high availability chassis cluster, the host logs updated in the system log messages might not show the correct time stamp. As a workaround, convert the UTC time stamp to local time zone. PR1394778

Platform and Infrastructure

- 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
- On NFX150 devices, the **request vmhost reboot in *minutes*** command with a delay specified in minutes reboots the device immediately. PR1406018
- On NFX250 devices, when you issue the **request support information** command, the configuration and counter data are missing for JDM. PR1413674

SEE ALSO

New and Changed Features | **181**

Changes in Behavior and Syntax | **182**

Known Behavior | **183**

Resolved Issues | **187**

Documentation Updates | **188**

Migration, Upgrade, and Downgrade Instructions | **188**

Product Compatibility | **191**

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in Junos OS 18.4R2 for NFX 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:18.4R2

Interfaces

- On NFX250 devices, an SFP-T interface does not become active when it is plugged into a ge-12/0/0 or a ge-13/0/0 interface. PR1404756

Platform and Infrastructure

- JDM depends on the libvirtd deamon to manage the guest VM through cli. The libvirtd daemon was stuck and vjunos VM start up failed, which resulted in in-band connectivity issues, the guest VM could not start, and the console was hung. PR1314945
- Software upgrade does not delete all images from a previous installation. This occupies about 1GB of storage per upgrade and leads to depletion of storage after several upgrades. PR1408061
- The **NFX3/ACX5448:LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified** message is displayed when you commit the configuration on config prompt. As a workaround to exclude this from messages or syslogs, run the **set system syslog user * match "!(LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified)** and commit. PR1376665
- On NFX250 devices, the **request-load-configuration** command output from device does not match with 18.4 yang. PR1416106
- With VNF running when MTU is configured, the KVM crashes and VNF goes down. PR1417103

Resolved Issues:18.4R1

There are no fixed issues in Junos OS Release 18.4R1 for NFX Series.

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 documentation for NFX 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 for the NFX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

Upgrade and Downgrade Support Policy for Junos OS Releases

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

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

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

Basic Procedure for Upgrading to Release 18.4

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 18.4R2:

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 by 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 | **181**

Changes in Behavior and Syntax | **182**

Known Behavior | **183**

Known Issues | **185**

Resolved Issues | **187**

Documentation Updates | **188**

Product Compatibility | **191**

Product Compatibility

IN THIS SECTION

Hardware Compatibility

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

To determine the features supported on NFX Series devices in this release, use the Juniper Networks Feature Explorer, a Web-based application that helps you 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.

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 platforms.

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 3 on page 191.

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 platform:

Table 3: 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

Table 3: Software Compatibility Details with vSRX and Cloud CPE Solution *(continued)*

NFX250 Junos OS Release	vSRX	Cloud CPE Solution
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-D45.3	15.1X49-D61	Not applicable
17.2R1	15.1X49-D78.3	Not applicable
17.3R1	15.1X49-D78.3	Not applicable
17.4R1	15.1X49-D78.3	Not applicable
15.1X53-D471	15.1X49-D143	Not applicable
18.1R1	18.1R1	Not applicable
18.1R2	18.1R2	Not applicable
18.1R3	18.1R3	Not applicable
18.2R1	18.2R1	Not applicable
18.3R1	18.3R1	Not applicable
18.4R1	18.4R1	Not applicable

SEE ALSO

New and Changed Features | **181**

Changes in Behavior and Syntax | **182**

Known Behavior | **183**

Known Issues | **185**

Junos OS Release Notes for PTX Series Packet Transport Routers

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 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 18.4R2 New and Changed Features

There are no new features or enhancements to existing features for PTX Series in Junos OS Release 18.4R2.

Release 18.4R1 New and Changed Features

Hardware

- **New fixed-configuration packet transport router (PTX Series)**—Starting in Junos OS Release 18.4R1, the PTX10001-20C is a new fixed-configuration Macsec-enabled LSR core router. It features a compact, 1U form factor that is easy to deploy in space-constrained Internet exchange locations, remote central offices, and embedded peering points throughout the network. The PTX10001 has 20 QSFP28 ports, and you can add 16 more QSFP28 ports with the optional JNP10001-16C-PIC expansion module. The 36 QSFP28 ports can be configured as 10 Gbps, 40 Gbps, or 100 Gbps. The ports handle up to 3.6 Tbps of throughput and 2 Bpps of forwarding capacity.

  See PTX10001 Hardware Guide.

Authentication, Authorization and Accounting (AAA) (RADIUS)

- **Support for password change policy enhancement (PTX Series)**—Starting in Junos OS Release 18.4R1, the Junos OS password change policy for local user accounts is enhanced to comply with additional password policies. As part of the policy improvement, you can configure the following:
  - **maximum-lifetime-value**—The maximum duration of a password. The password expires after the maximum is reached.
  - **minimum-lifetime-value**—The minimum duration of a password. You cannot change the password until the minimum duration is reached.

  [See password.]

Class of Service (CoS)

- **Support for classifying Layer 2 frames based on Layer 3 information (PTX Series)**—Starting in Junos OS Release 18.4R1, PTX Series devices support classifying Layer 2 frames based on Layer 3 fields. You can match on DSCP bits in IPv4 packets (classifier type **dscp**), TOS bits in IPv6 packets (classifier type **dscp-ipv6**), EXP bits in MPLS frames (classifier type **exp**), and PCP bits in IEEE 802.1 frames (classifier type **ieee-802.1**). To do this, define classifiers as normal at the **[edit class-of-service classifiers *classifier-type classifier-name*]** hierarchy level and then apply the classifiers to a Layer 2 (**family ethernet-switching**) interface at the **[edit class-of-services interfaces *interface-name* unit 0]** hierarchy level.

  [See classifiers (Definition).]

- **Support for class of service (CoS) on PTX10001-20C routers**—Starting in Junos OS Release 18.4R1, PTX10001-20C routers support class-of-service (CoS) functionality.

  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 CoS Features and Limitations on PTX Series Routers.]

Forwarding and Sampling

- **Support for activating or deactivating static routes on the basis of RPM test results (PTX Series)** —Starting in Junos OS 18.4R1, you can use RPM probes to detect link status, and change the preferred-route state on the basis of the probe results. Tracked routes can be IPv4 or IPv6, and support a single IPv4 or IPv6 next hop. For example, RPM probes can be sent to an IP address to determine if the link is up, and if so, take the action of installing a static route in the route table. RPM-tracked routes are installed with preference 1 and thus are preferred over any existing static routes for the same prefix.

  [See Configuring RPM Probes , rpm-tracking, and show route rpm-tracking.]

Interfaces and Chassis

- **LACP hold-up timer configuration support on LAG interfaces (PTX Series)**—You can configure an LACP hold-up timer value for LAG interfaces to prevent excessive flapping of a child (member) link of a LAG interface due to transport layer issues.

  Because of transport layer issues, a link can be physically up and still cause LACP state-machine flapping. LACP state-machine flapping, which can adversely affect traffic on the LAG interface. With the hold-up timer 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 hold-up timer, 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.]

Junos Telemetry Interface

- **Enhanced IS-IS sensor support for Junos Telemetry Interface (JTI) (MX960, MX2020, PTX5000, PTX1000, and PTX10000)**—Starting with Junos OS Release 18.4R1, JTI supports OpenConfig Version v0.3.3 (from v0.2.1) for resource paths related to IS-IS Link State Database (LSDB) streaming. The difference between the two versions results in changes, additions, deletions, or nonsupport for leaf devices related to the following IS-IS Type Length Value (TLV) parameters and IS-IS areas:
  - TLV 135: extended-ipv4-reachability
  - TLV 236: ipv6-reachability
  - TLV 22: extended-is-reachability

- TLV 242: router-capabilities
- IS-IS Interface Attributes
- IS-IS Adjacency Attributes

To provision the sensor to export data through gRPC streaming, use the **telemetry Subscribe** RPC to specify telemetry parameters. Streaming telemetry data through gRPC also requires the OpenConfig and Network Agent packages, both of which are bundled into the Junos OS image in a default package named junos-openconfig.

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

- **Support for NTF agent (MX240, MX480, MX960, MX2010, MX2020, vMX, PTX1000, PTX5000, and PTX10000)**—Junos OS exposes telemetry data over gRPC and UDP as part of the Junos Telemetry Interface (JTI). One way to stream JTI data into your existing telemetry and analytics infrastructure requires managing an external entity to convert the data into a compatible format. Starting in Junos OS Release 18.4R1, the NTF agent feature provides an on-box solution that enables you to configure and customize to which endpoint (such as IPFIX and Kafka) the JTI data is delivered and in which format (such as AVRO, JSON, and MessagePack) the data is encoded.

  [See NTF Agent Overview.]

- **Expanded ON_CHANGE support for Junos Telemetry Interface (JTI) (MX960, MX2010, MX2020, PTX5000, PTX1000, and PTX10000)**—Starting in Junos OS Release 18.4R1, OpenConfig support through gRPC and JTI is extended to support additional ON_CHANGE sensors.

  Periodical streaming of OpenConfig operational states and counters collects information at regular intervals. ON_CHANGE support streams operational states as events (only when there is a change), and is preferred over periodic streaming for time-sensitive missions.

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

  - **/components/component**
  - **/components/component/name/**
  - **/components/component/state/type**
  - **/components/component/state/id**
  - **/components/component/state/description**
  - **/components/component/state/serial-no**
  - **/components/component/state/part-no**

  ON_CHANGE notification will be supported on all the hardware components displayed in the Junos OS CLI operational mode command **show chassis hardware**.

To provision a sensor to export data through gRPC, use the **telemetrySubscribe** RPC to specify telemetry parameters. To enable ON_CHANGE support, configure the sample frequency in the subscription as zero.

[See Guidelines for gRPC Sensors (Junos Telemetry Interface) and show chassis hardware.]

Layer 2 Features

- **Support for Layer 2 and Layer 3 forwarding across VLANs (PTX1000, PTX10008, and PTX10016)**—Starting in Junos OS 18.4R1, PTX Series devices support Layer 2 and Layer 3 forwarding across VLANs. Layer 3 forwarding across VLANs by using Integrated Routing and Bridging (IRB) interface. To provide Layer 3 forwarding across VLANs, you need to create layer 3 logical interface on IRB physical interface and associate it with the VLAN.

  PTX Series routers support enabling IS-IS and OSPF protocols at the IRB level and also support iBGP and eBGP on the IRB interface. You can apply firewall filter and policer on the IRB interface.

  [See Layer 2 Learning and Forwarding for VLANs Overview.]

- **Support for port mirroring (PTX10001)**—Starting in Junos OS Release 18.4R1, the PTX10001 supports firewall filter-based port mirroring for the IPv4 address family on the ingress interface.

  [See Configuring Port Mirroring on M, T MX, and PTX Series Routers.]

Layer 3 Features

- **Support for BFD on PTX10001-20C Packet Transport Router**—Starting in Junos OS Release 18.4R1, PTX10001-20C routers support Bidirectional Forwarding Detection (BFD) in centralized mode for clients operating under Layer 3 protocols such as OSPF, IS-IS, and BGP. BFD support is not extended to micro-BFD, IPv6, PIM, tunnel interfaces, or MPLS. [See bfd command.]

- **Support for ECMP on Layer 3 and MPLS routes on PTX10001-20C Packet Transport Router**—Starting in Junos OS Release 18.4R1, PTX10001-20C routers support equal-cost multipath (ECMP) load balancing for IPv4 and MPLS routes.

- **Support for Layer 3 unicast features on PTX10001-20C Packet Transport Router** —Starting in Junos OS Release 18.4R1, PTX10001-20C routers support the following Layer 3 forwarding features for unicast IPv4 traffic:

  - ICMPv4 messages (MTU exceeded, TTL expiry, host unreachable, IP redirect)
  - ICMPv4 host and longest prefix match (LPM) routing
  - IP packet exceptions (TTL error and IP-option)
  - IPv4 fragmentation
  - IPv4 ping and traceroute
  - Layer 3 protocols, such as:
    - OSPF
    - IS-IS with Bidirectional Forwarding Detection (BFD)

- BGP
- MTU check per port
- Virtual router (VRF-lite)

MPLS

- **MPLS support (PTX10001-20C)**—Starting with Junos OS Release 18.4R1, MPLS is supported on the PTX10001-20C router. The following features are supported:
  - Label Switching Routers (LSRs)
  - LDP and RSVP MPLS routing protocols
  - IS-IS interior gateway protocol (IGP) traffic engineering
  - Object access method, including ping and Bidirectional Forwarding Detection (BFD)
  - Fast reroute (FRR) MPLS local protection. Both one-to-one local protection and many-to-one local protection are supported.

  This feature was previously supported in an "X" release of Junos OS. [See MPLS Overview.]
- **MPLS-TE Fast Reroute Link Protection (PTX10001-20C)**— Starting with Junos OS Release 18.4R1, you can enable fast reroute (FRR) to automatically reroute traffic on MPLS traffic engineering (TE) LSPs if a node or link in an LSP fails, thus reducing the loss of packets traveling over the LSP. When you enable fast reroute, detours are precomputed and pre-established along the LSP. In case of a network failure on the current LSP path, traffic is quickly routed to one of the detours. Fast reroute protects traffic against any single point of failure between the ingress and egress routers.

  This feature was previously supported in an "X" release of Junos OS. [See Fast Reroute Overview.]

Network Management and Monitoring

- **sFlow functionality introduced on PTX1000 and PTX10000**—Starting in Junos OS Release 18.4R1, 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.]

Port Security

- **Media Access Control Security (MACsec) support (PTX10001-20C routers)**—Starting in Junos OS Release 18.4R1, MACsec is supported on all twenty interfaces on the PTX10001-20C router and all sixteen interfaces on the TIC1 module. 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).]

- **Dynamic Host Configuration Protocol (DHCP) relay (PTX10001-20C routers)**—Starting in Junos OS Release 18.4R1, DHCP relay is supported on PTX10001-20C routers.

  [See Extended DHCP Relay Agent.]

Routing Protocols

- **Support for 64 add-path BGP routes (PTX Series)**—Starting in Junos OS Release 18.4R1, support is extended to 64 add-path BGP routes. Currently Junos OS supports six add-path routes and BGP can advertise upto 20 add-path routes through policy configuration. This feature allows BGP to advertise 64 add-path routes and a second best ECMP path as a backup in addition to the multiple ECMP paths.

  To advertise all add-paths up to 64 add-paths or only equal-cost paths, include the **path-selection-mode** statement at the **[edit protocols bgp group *group-name* family *name* addpath send]** hierarchy level. You cannot enable both **multipath** and **path-selection-mode** at the same time.

  To advertise a second best ECMP path as a backup path in addition to the multiple ECMP paths include the **include-backup-path *bacup_path_name*** statement at the **[edit protocols bgp group *group-name* family *name* addpath send]]** hierarchy level.

  [See add-path.]

  [See include-backup-path.]

- **Support for BGP flowspec redirect to IP (PTX Series)**—Starting in Junos OS Release 18.4R1, BGP flow specification as described in BGP Flow-Spec Internet draft draft-ietf-idr-flowspec-redirect-ip-02.txt, *Redirect to IP Action* is supported. Redirect to IP action uses extended BGP community to provide traffic filtering options for DDoS mitigation in service provider networks. Legacy flow specification, as specified in the Internet draft draft-ietf-idr-flowspec-redirect-ip-00.txt, *BGP Flow-Spec Extended Community for Traffic Redirect to IP Next Hop*, redirect to IP uses the BGP nexthop attribute to support interoperability of devices. Junos OS advertises redirect to IP flow specification action using the extended community by default. Redirect to IP action allows you to divert matching flow specification traffic to a globally reachable address. This feature is required to support service chaining in virtual service control gateway (vSCG).

  To configure a static IPv4 flow specification route, include the **redirect *ipv4-address*** statement at the **[edit routing-options flow route then]** hierarchy level in the configuration.

  To configure a static IPv6 specification route, include the **redirect *ipv6-address*** statement at the **[edit routing-options flow route then]** hierarchy level in the configuration.

  To configure legacy flow specification include **legacy-redirect-ip-action** at the **[edit group *bgp-group* neighbor *bgp neighbor* family inet flow]** hierarchy level.

  To configure BGP to use VRF.inet.0 table to resolve VRF flow specification routes, include **secondary-independent-resolution** statement at the **[edit protocols bgp neighbor family flow]** hierarchy level.

  [See legacy-redirect-ip-action.]

[See Configuring BGP Flow Specification Action Redirect to IP to Filter DDoS Traffic.]

Security

- **Support for Ingress Firewall Filters (PTX10001-20C)**—Starting with Junos OS Release 18.4R1, you can configure firewall rules to filter incoming network traffic based on a series of user-defined rules. You can specify whether to accept, permit, deny, or forward packets before it enters an interface. If a packet is accepted, you can also configure additional actions to perform on the packet, such as class-of-service (CoS) marking (grouping similar types of traffic together and treating each type of traffic as a class with its own level of service priority) and traffic policing (controlling the maximum rate of traffic sent or received). Only ingress firewall filters are supported. You configure firewall filters under the **[edit firewall]** hierarchy level. This feature was previously supported in an "X" release of Junos OS.

  [See Firewall Filters Overview.]

Services Applications

- **Support for IPv4 and IPv6 inline active flow monitoring (PTX10002-60C router)**—Starting in Junos OS Release 18.4R1 on PTX10002-60C routers, you can perform inline active flow monitoring for IPv4 and IPv6 traffic. Both IPFIX and version 9 templates are supported.

  [See Configuring Inline Active Flow Monitoring on PTX Series Routers.]

System Management

- **Copy files between the Junos VM and Linux host (PTX10008)**—In Junos OS Release 18.4R1, two commands are introduced on the Enhanced Automation variant of Junos OS for PTX10008 routers: **request vmhost copy jnode-to-vjunos** and **request vmhost copy vjunos-to-jnode**. These commands enable you to copy files from the Linux host to the Junos VM and vice versa.

  [See request vmhost copy jnode-to-vjunos and request vmhost copy vjunos-to-jnode.]

VPN

- **Support to control traceroute over Layer 3 VPN (PTX Series)**—Starting in Junos OS Release 18.4R1, in a Layer 3 VPN topology with **vrf-table-label** configured and multiple customer edge (CE) routers configured in the same VPN routing and forwarding (VRF) routing instance, when traceroute is performed to a remote provider edge (PE) router for a CE-facing network, the ICMP time exceeded packet determines the correct IP address as the source address.

  To control the traceroute over Layer 3 VPN topology with vrf-table-label configured and multiple CE routers configured in the same VRF, you can configure **allow-l3vpn-traceroute-src-select** at the**[edit system]** hierarchy level that determines the correct IP source address by reviewing the destination routing instance and destination IP address.

  [See allow-l3vpn-traceroute-src-select.]

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 main release and the maintenance releases for the PTX Series.

Release 18.4R2-S1 Changes in Behavior and Syntax

Software Defined Networking (SDN)

- **Increase in the maximum value of delegation-cleanup-timeout (PTX Series)**—You can now configure a maximum of *2147483647* seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

  With the increase in maximum value of **delegation-cleanup-timeout** from *600* to *2147483647* seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

  [See delegation-cleanup-timeout.]

Release 18.4R2 Changes in Behavior and Syntax

Interfaces and Chassis

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (PTX Series)**—In Junos OS Release 18.4R2, 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 before an interface changes state from down to up. In earlier Junos OS releases, the LACP hold-up the information for all interfaces was in a single **<lacp-hold-up-information>** XML tag. Now, the hold-up information for each interface is displayed in a separate **<lacp-hold-up-information>** XML tag.

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

- **Change in error severity (PTX10016)**—Starting in Junos OS Release 18.4R2, on PTX10016 routers, the severity of the FPC error, shown in the syslog as **PE Chip::FATAL ERROR!! from PE2[2]: RT: Clear Fatal if it is detected LLMEM Error MEM:llmem, MEMTYPE: 1**, is changed from fatal to non-fatal (or minor). In case of this error, only a message is displayed for information purpose. To view the error details, you can use the show commands **show chassis fpc errors** and **show chassis errors active**.

  [See show chassis fpc errors.]

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 18.4R2, under rare circumstances, the host interface may stop sending packets and the connections to and from the peer might fail if an outbound firewall filter is configured with the **syslog** and **accept** actions. This condition applies to IPv4 and IPv6 traffic families. We recommends that you do not use the **syslog** and **accept** actions in the output filter for these systems.

Here's a sample configuration (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
```

[See PR 1354580.]

Release 18.4R1 Changes in Behavior and Syntax

Interfaces and Chassis

- **New option to configure IP address to be used when the Routing Engine is the current master**—Starting in Junos OS Release 18.4R1, 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 releases before Junos OS Release 18.4R1, the same IP address is applied on the master and backup Routing Engines, resulting in configuration issues.
- **Support for creating Layer 2 logical interface independently (PTX Series)**—In Junos OS Releases 18.4R1, 18.4R2, and later, PTX Series routers support creating layer 2 logical interface independent of layer 2 routing instance type. That is, you can configure and commit the layer 2 logical interfaces separately and add the interface to bridge-domain or Ethernet VPN (EVPN) routing instance separately. Note that the layer 2 logical interfaces works fine only when the interface is added to bridge domain or EVPN routing instance.

  In the earlier Junos OS releases, when an layer 2 logical interface configuration (units with encapsulation vlan-bridge configuration) is used, then the logical interface must be added as part of a bridge-domain or EVPN routing instance for the commit to succeed.

Network Management and Monitoring

- **No chassis alarm when power consumption by an FPC exceeds 90% or 100% of the allocated power budget**—Starting in Junos OS Release 18.4R1, 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 18.4R1, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, the server must not return an RPC reply that encloses both an **<rpc-error>** element and an **<ok/>** element. If the operation is successful, but the server reply encloses one or more **<rpc-error>** elements of severity warning in addition to the **<ok/>** element, then 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 encloses both an **<rpc-error>** element of severity warning and an **<ok/>** element.
- **Deque Dry Interrupt error severity changed to fatal**—Starting in Junos OS Release 18.4R1, on PTX5000 routers, we have changed the severity of the error Deque Dry Interrupt (error code: 0x2100dd) from major to fatal. By default, this error disables the Packet Forwarding Engine on the FPC. You can use the **show chassis fpc errors** command to view the default or user-configured action that resulted from the error.

  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 1-888-31 4-JTAC (within the United States) or 1-408-7 45-9500 (from outside the United States).

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 18.4R2 for the PTX Series.

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

Interfaces and Chassis

- On PTX10001-20C routers, the **show interfaces** command might display different values for the input and output packets per second (pps) for host-bound packets.

General Routing

- When an FPC goes offline or restarts, a source FPC sends traffic to a destination FPC. The following error messages are seen and a corresponding alarm is set on the destination FPC. Specific to PTX10000, the transient alarm gets set when this condition occurs. The alarm clears later because the source FPC goes offline. **Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210613), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000010: Grant spray drop due to unspray-able condition error Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Error (0x210614), module: PE Chip, type: Apr 09 10:31:24 [TRACE] [asta] Apr 9 10:19:59 asta fpc4 Cmerror Op Set: PE Chip: PE1[1]: FO:core intr: 0x00000008: Request spray drop due to unspray-able condition error** PR1268678
- The Routing Engine boots from the secondary disk when you:

  Press the reset button, on the RCB front panel, while Routing Engine is booting up but before Junos 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.
  - Reboot and the system hangs before Junos is up PR1344342
- Recommendation: The ingress interface and the mirror interface should have the same MTU or you must set a higher MTU on the mirror interface than on the ingress interface. PR1372321
- Currently, PTX1000-M20C supports 128,000 transit LSPs; however, in a failover scenario, Argus can support a maximum of 192,000 LSPs, which means 64,000 backup LSPs are active. In a failover scenario and MBB case, 256,000 LSPs are required, but the ASIC can handle a maximum of 192,000 after optimization, so there is a limitation with backup LSPs. PR1375780
- PTX1000 and MX Series sFlow sampling output has different VLAN priority in extended switch data fields with the same dual-tag configuration when egress sampling is configured, the difference is due to

the sequence in which sampling and mac-rewrite happens. In MX Series, MAC rewrite occurs after sampling, and in the case of PTX Series sampling, happens after MAC rewrite. PR1387468

- **set interfaces *interface-name* gigether-options fec <fec74/fec91/none>** configuration is not supported on Argus platform. PR1388140

User Interface and Configuration

- **Auto-complete caution for QFX10002-60c and PTX10002-60c personalities**—Starting in Junos OS Release 18.4R1, for QFX10002-60c and PTX10002-60c personalities, do not use auto-complete to display the list of arguments for the **request system software delete** command. You must look for the package name using the **show system software** command and then explicitly type the software package name in the **request system software delete** command.

  [See request system software delete].

SEE ALSO

- New and Changed Features | **193**
- Changes in Behavior and Syntax | **201**
- Known Issues | **206**
- Resolved Issues | **211**
- Documentation Updates | **216**
- Migration, Upgrade, and Downgrade Instructions | **216**
- Product Compatibility | **220**

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.4R2 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.

Class of Service (CoS)

- Core files are generated when ports are channelized and de-channelized repeatedly, without delay. PR1370781

General Routing

- Control packets might get dropped when the Packet Forwarding Engine experiences heavy congestion. PR1163759
- In a rare race condition, multiple interrupts are not handled properly on MX Series device with MPC7E, MPC8E, or MPC9E and on PTX Series devices with FPC3-PTX-U2 or FPC3-PTX-U3, which could lead to 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
- On the third-generation PTX Series FPCs (PTX3000 or PTX5000 FPC3, PTX1000) if the **protocols mpls no-propagate-ttl** statement is configured, the MPLS TTL field can be reset to 255 in the packets where a label swap operation is performed. PR1287473
- On a PTX Series PIC with the CFP2-DCO-T-WDM transceiver installed, after repeated configuration rollback, the link sometimes takes a long time to come up. PR1301462
- On a PTX Series router with a third-generation FPC, an error message is displayed when the FPC comes online or goes offline. PR1322491
- On a PTX Series router, this error message is seen in the Packet Forwarding Engine syslog on every FPC reboot: **SCHED: Thread 57 (CMSNGFPC) ran for 2002 ms without yielding or [...LOG: Emergency] SCHED: Thread 50 (CMSNGFPC) aborted, hogged 8899 ms**. There is no functional impact, so that the error can be ignored. PR1343256
- On QFX10000 switches or on PTX Series routers, NETCONF over SSH traffic through TCP port 830 might hit the unclassified host path queue. This can result in DDoS violations in the unclassified queue. PR1345744

- PTX3000 reports Chip to Chip Link (CCL) CRC errors while the FPC3-SFF-PTX-1X is brought offline using a CLI command or by pressing the offline button. The syslog error is generated by an FPC just before it goes offline, so there is no detectable traffic loss. 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
- When you commit an aggregated Ethernet configuration, harmless errors are seen. PR1365355
- When the TIC goes offline and then comes back online, MPLS bidirectional traffic flow might stop working. PR1367920
- Unsuccessful connection attempts are not logged on the backup SPMB. PR1369731
- 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 in to the UKERN running on the FPC, which drops these packets. The messages seen here are displayed just before the packets are dropped. These error messages are harmless and do not disrupt the working of any feature. PR1374372
- In certain scenarios where flows are sampled through aggregated bundles when J-flow 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
- The DHCP relay functionality does not work on PTX10001-20C devices. DHCP relay functionality: The DHCP requests and the DHCP offers are snooped by the box, the snooping happens through firewall, which snoops all the DHCP packets entering the default route table, and all the offers and requests are punted unto the host or control-plane. When a DHCP client sends the DHCP request, it gets intercepted by the filter block and punted up to the control plane. Upon receiving this packet, the control plane unicasts (relays) this packet to DHCP servers. The DHCP server responds with a DHCP offer, which again gets intercepted by the firewall block and punted up. Upon receiving the DHCP offer, the control plane broadcasts this DHCP offer to the client's VLAN and eventually the client receives the DHCP offer. PR1407476
- When a 100g QSFP is inserted into FPC on PTX Series routers, 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
- The **rx_power** value streamed to the telemetry server is the raw value (in mW) returned directly from the transceiver driver. The Junos OS CLI value has been transformed in the transportd process into different units: (Rx input total power(0.01dBm). PR1411023
- This issue is specific to PTX10002. During normal operation, if the **chassis-control** process restarts, certain ASICs are not initialized. This causes packet drops on the output queue. PR1414434

- In PTX3000 system, only if the IPLC card is present in the system, and when GRES is performed, we will observe IPLC crash during the GRES operation. There is no impact on other line cards in the system. If there is no IPLC card in the system, there is no impact during the GRES. PR1415145
- On an MX Series, performing the **show forwarding-options load-balance ...** command might cause Packet Forwarding Edge wedging after a certain number of attempts (less than 200 in test) if the **destination-address** option of the command matches the default route with the **discard** action. This is because a defect code causing internal flow errors is involved in that scenario. PR1422464
- On routers and switches running Junos OS, with Link Aggregation Control Protocol (LACP) enabled, deactivating a remote aggregate Ethernet member link makes the local member link move to LACP Detached state. The detached link is then 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 are enabled in Packet Forwarding Engine. Aggregated Ethernet-Forwarding table irrespective of LACP states, which result in traffic drop. PR1423707
- When one of the PEMs is not present or not powered, an active alarm should be flagged and a syslog, indicating the same should be generated. But due to the defect, it does not occur. PR1439198

Infrastructure

- Junos OS packages might have been 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 the upgrade. This is because of the old version of **/var/db/cfm.db**. PR1281073

MPLS

- On devices running Junos OS, with transit chaining mode enabled, if RSVP link/node protection is enabled and **sensor-based-stats** is used, a single-hop bypass label-switched path (LSP) next hop might not be installed in forwarding information base (FIB) even it is in routing information base (RIB). Thus the single-hop bypass LSP will fail to forward traffic when needed. PR1401152

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 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

SEE ALSO

New and Changed Features | **193**

Changes in Behavior and Syntax | **201**

Known Behavior | **204**

Resolved Issues | **211**

Documentation Updates | **216**

Migration, Upgrade, and Downgrade Instructions | **216**

Product Compatibility | **220**

Resolved Issues

IN THIS SECTION

This section lists the issues fixed in the Junos OS Release 18.4R2 for the PTX Series.

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

General Routing

- On PTX Series, multicast traffic packet drop of more than 50 percent is seen when a first-generation or second-generation FPC is used in the same chassis with a third-generation FPC. PR1339481
- Disable reporting of correctable single-bit errors on Hybrid Memory Cube (HMC) and prevent a major alarm. PR1384435
- Packet drop might be seen in lower-priority queues on PTX Series routers or on the QFX10000 line of switches. PR1385454
- The **show chassis fpc** command on PTX1000 routers and the PTX10000 line of routers shows incorrect buffer memory utilization. PR1397612
- High jsd or na-grpcd CPU usage might be seen even if JET or JTI is not used. PR1398398
- CPU overuse might be observed on PTX Series routers or on the QFX10000 line of switches. PR1399369
- The DHCPv6 relay-reply packet might be dropped by the DHCP relay. PR1399683

- Only one Packet Forwarding Engine could be disabled on an FPC with multiple Packet Forwarding Engines in an error or wedge condition. PR1400716
- The TCP connection between ppmd and ppman might be dropped because of a kernel issue. PR1401507
- Log message **JAM HW data base open failed for ptx5kpic_3x400ge-cfp8** during commit. PR1403071
- Incorrect mem stat message is seen in FPC logs of PTX Type 1 FPC. PR1404088
- RPT TPTX REGRESSIONS: While checking ethernet-switch verification ethernet-switch statistics is not in expected range. PR1404365
- On a PTX3000, FPCs are not able to come online for tens of minutes after a reboot of the chassis. PR1404611
- ZTP upgrade might fail if there are more than one 10-Gigabit Ethernet interfaces connected to the DHCP server. PR1404832
- On PTX3000 or PTX5000, the backup CB's chassis environment status is always**Testing** you remove and reinsert the backup CB. PR1405181
- 100-gigabit SR4 optics with part number 740-061405 should be displayed as **QSFP-100G-SR4-T2**. PR1405399
- No chassis alarm is raised on PTX1000 when the PEM is removed or power lost to PEM. PR1405430
- Layer2 VPN might flap repeatedly when the link between the PE device and CE device is coming up. PR1407345
- The Packet Forwarding Engine might get disabled unexpectedly due to a auto correctable non-fatal hardware error on PTX Series routers or QFX10002, QFX10008, or QFX10016. PR1408012
- openconfig-network-instance:network-instances support for IS-IS must be hidden unless supported. PR1408151
- PTX Inline J-flow: FPC went offline when sampling rate was changed at runtime to 80,000; dcpfe core file was also generated. PR1409502
- The CPU might be overused by jsd process in JET scenario. PR1409639
- Hostname is not updated at the FPC shell after a system configuration change on the CLI. PR1412318
- Junos PCC might reject PCUpdate/PCCreate message if there is metric type other than type 2. PR1412659
- The Layer 2 circuit egress PE device might drop the traffic in a FAT+CW-enabled Layer 2 circuit scenario when another FAT+CW enabled Layer 2 circuit PW flaps. PR1415614
- Traffic loss could be seen for the duration of the hold-time down timer when an interface, with the hold-time down timer configured, flaps. with hold-time down timer configured. PR1418425
- RX alarms are not set according to the threshold value configured for the DCO Tunable Optics. PR1419204
- An interface might go to down state on a QFX10000 or PTX10000 platform. PR1421075

- Virtual Chassis might become unstable and fxpc core files might be generated when there are a lot of configured filter entries. PR1422132
- 4x10G interfaces on the third-generation FPCs on PX3000 or PTX5000 might not come up after frequently flap ping for a long of time. PR1422535
- While committing a huge configuration, the user sees the error **error: mustd trace init failed**.PR1423229
- A Specific interface on the 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. PR1425003

Interfaces and Chassis

- The syslog message **/kernel: %KERN-3: pointchange for flag 04000000 not supported on IFD aex** upon LFM related configuration commit on aggregated Ethernet interfaces. PR1423586
- Some ports on PTX Series routers might remain down after rebooting the FPC or the device is rebooted at the 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 an indirect next-hop is deleted. PR1398876
- LDP routes might flap if committing any configuration changes. PR1416032
- Bypass dynamic RSVP LSP tears down too soon when being used for protecting an LDP LSP with **dynamic-rsvp-lsp** statement. PR1425824

Platform and Infrastructure

- Some files are missing during log archiving. PR1405903

Routing Protocols

- Rpd core files are seen on the backup Routing Engine during neighborship flapping when the **authentication-key** option with a size larger than 20 characters is used. PR1394082
- Syslog message is seen whenever the prefix SID coincides with the node SID. PR1403729

- An rpd memory leak might be seen in an IS-IS segment routing scenario. PR1404134
- Dynamic routing protocol flapping with VM host Routing Engine switchover on NG-RE. PR1415077
- Rpd might crash with ospf overload configuration. PR1429765

Infrastructure

- The FPC might go down on some VM-host-based PTX Series or QFX Series devices. PR1367477

Interfaces and Chassis

- **PE Chip:pe0[0]: IPW: oversize_drop error** causes a major error on FPC. PR1375030

MPLS

- In Junos OS Release 18.2X75, IPv6 routes are dead in mpls.0 table S=0 leads to traffic loss in v6-indirect next-hop stitching. PR1355878
- LSP with **auto-bandwidth** enabled goes down as a result of an HMC error. PR1374102
- Bypass LSP is taking the same SRLG colored path. PR1387497

Platform and Infrastructure

- On a PTX1000, upgrade from Junos OS Release 16.1X65-D45 to Junos OS Release 17.3-20170721 fails frequently when sampling is enabled. PR1296533
- Repeated log messages **%PFE-3 fpcX expr_nh_index_tree_ifl_get** and **expr_nh_index_tree_ipaddr_get** are observed when the sampling packet is discarded with the log (or syslog) statement configured under the firewall filter. PR1304022
- The status LED on the chassis remains unlit on the QFX10002-60C. PR1332991
- The traffic-class-count values in a filter configured with the **scale-optimized** statement, are not incrementing. PR1334580
- Packet might be dropped by RPF during a Routing Engine switchover. PR1354285
- The host interface might stop sending packets on a PTX Series router with FPC3 or PTX1000 when you use an outbound firewall filter with **syslog** option. PR1354580
- PTX1000-M20C: FRR link-protection convergence time. PR1355953
- Traffic is still forwarded through the member link of an aggregated Ethernet bundle interface even with **Link-Layer-Down** flag set. PR1365263

- JSA10899 2018-10 Security Bulletin: Junos OS: Next-hop index allocation failed: private index space exhausted as a result of incoming ARP requests to the management interface (CVE-2018-0063). PR1360039
- The 'Normal discards' Packet Forwarding Engine statistics traffic counter might increase at a higher rate when Inline-Jflow or sFlow is enabled. PR1368208
- slu.l2_domain_lookup_failure traps might be observed when using sampling on FPC-P1/FPC-P2. PR1368381
- The IPLC card might take a long time to come up. PR1368637
- The 'commit or commit check operation' might fail because of the error **cannot have lsp-cleanup-timer without lsp-provisioning**.PR1368992
- On PTX10001 and ACX6360, 100G-LR4 optics and 100G-ER4 optics are not supported. PR1371590
- Packets might be dropped after a filter is deleted from an interface. PR1372957
- Inline BFD keeps flapping when inline samping is configured. PR1376509
- Traffic might be dropped on third-generation FPCs on PTX Series routers. PR1378392
- Layer 3 VPN traffic might be dropped because one core-facing interface is down. PR1380783
- BFD sessions bounced FPCs that have not been taken offline. PR1383703
- Packet Forwarding Engine-based local repair does not happen for IP routes pointing to a unilist of composites with Indirect next hops. PR1383965
- CPSM daemon memory leak is observed ion VM host. PR1387903
- BFD flaps are seen on PTX or QFX10000 platforms with inline BFD. PR1389569
- Forwarding issue on mixed link-speed aggregated Ethernet interface after FPC reloads. PR1390417
- High jsd or na-grpcd CPU usage might be seen even when JET or JTI is not used. PR1398398

SEE ALSO

New and Changed Features | **193**

Changes in Behavior and Syntax | **201**

Known Behavior | **204**

Known Issues | **206**

Documentation Updates | **216**

Migration, Upgrade, and Downgrade Instructions | **216**

Product Compatibility | **220**

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 documentation for PTX Series.

SEE ALSO

- New and Changed Features | 193
- Changes in Behavior and Syntax | 201
- Known Behavior | 204
- Known Issues | 206
- Resolved Issues | 211
- Migration, Upgrade, and Downgrade Instructions | 216
- Product Compatibility | 220

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 PTX Series. Upgrading or downgrading Junos OS might take several hours, depending on the size and configuration of the network.

Basic Procedure for Upgrading to Release 18.4

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 18.4R2:

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 by using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.

7. Review and accept the End User License Agreement.

8. Download the software to a local host.

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

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

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

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-18.4R2.9.tgz
```

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

```
user@host> request system software add validate reboot
source/junos-install-ptx-x86-64-18.4R2.9-limited.tgz
```

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

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

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

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

Rebooting occurs only if the upgrade is successful.

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

NOTE: After you install a Junos OS Release 18.4 **jinstall** package, you cannot return to the previously installed software by issuing the **request system software rollback** command. Instead, you must issue the **request system software add validate** command and specify the **jinstall** package that corresponds to the previously installed software.

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

Upgrade and Downgrade Support Policy for Junos OS Releases

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

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from Junos OS Release 17.1 to Release 17.2 or from Junos OS Release 17.1 to Release 17.3. However, you cannot upgrade directly from a non-EEOL release that is more than three releases ahead or behind.

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

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

Upgrading a Router with Redundant Routing Engines

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

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

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

SEE ALSO

- New and Changed Features | **193**
- Changes in Behavior and Syntax | **201**
- Known Behavior | **204**
- Known Issues | **206**
- Resolved Issues | **211**
- Documentation Updates | **216**
- Product Compatibility | **220**

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | **221**

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

Hardware Compatibility Tool

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

SEE ALSO

- New and Changed Features | **193**
- Changes in Behavior and Syntax | **201**
- Known Behavior | **204**
- Known Issues | **206**
- Resolved Issues | **211**
- Documentation Updates | **216**
- Migration, Upgrade, and Downgrade Instructions | **216**

Junos OS Release Notes for the QFX Series

IN THIS SECTION

- New and Changed Features | **222**
- Changes in Behavior and Syntax | **233**
- Known Behavior | **237**
- Known Issues | **240**
- Resolved Issues | **249**
- Documentation Updates | **262**

These release notes accompany Junos OS Release 18.4R2 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 for the QFX Series switches in Junos OS Release 18.4R2.

NOTE: The following QFX Series platforms are supported in Release 18.4R2: QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10008, and QFX10016. Junos on White Box is also supported in Release 18.4R1.

New and Changed Features: 18.4R2-S3

System Management

- **Additional support for Bidirectional Forwarding Detection (QFX5110, QFX5120, QFX5200, and QFX5210)**—Starting in Junos OS Release 18.4R2-S3, inline Bidirectional Forwarding Detection (BFD) is enabled by default.

  [See Understanding Bidirectional Forwarding Detection (BFD).]

New and Changed Features: 18.4R2

EVPNs

- **Layer 2 and Layer 3 VXLAN gateways (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 18.4R2, you can deploy EX4650 and QFX5120 switches as follows:

  - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
  - (QFX5120 switches only) As a Layer 2 VXLAN gateway in an Open vSwitch Database (OVSDB) overlay network

  VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

  [See Understanding VXLANs.]

- **EVPN control plane and VXLAN data plane support (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 18.4R2, EX4650 and QFX5120 switches support EVPN-VXLAN. By using a Layer 3 IP-based underlay network coupled with an EVPN-VXLAN overlay network, you can place endpoints anywhere in the network and remain connected to the same logical Layer 2 network.

  EVPN-VXLAN is commonly deployed over the following physical underlay architectures:

  - A two-layer IP fabric that includes spine devices (Layer 3 VXLAN gateways) and leaf devices (Layer 2 VXLAN gateways). You can deploy EX4650 and QFX5120 switches as spine or leaf devices in this fabric.
  - A one-layer IP fabric that includes leaf devices that function as both Layer 2 and Layer 3 VXLAN gateways. You can deploy EX4650 and QFX5120 switches as leaf nodes in this fabric.

  [See Understanding EVPN with VXLAN Data Encapsulation.]

- **EVPN pure type-5 route support (EX4650 and QFX5120 switches)**—Starting with Junos OS Release 18.4R2, you can configure pure type-5 routing in an EVPN-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 carries the

MAC address of the sending switch and provides 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 ip-prefix-routes.]

- **Selective multicast forwarding and SMET support in EVPN-VXLAN (QFX5110 and QFX5120 switches)**—Starting in Junos OS Release 18.4R2, Junos OS supports selective multicast Ethernet forwarding in an EVPN-VXLAN network. IGMP snooping enabled devices on a bridge domain monitor and selectively forward traffic from the access interface to the core. Devices that support selective multicast Ethernet forwarding do not send multicast traffic to all devices. Instead, they replicate and forward multicast traffic only to the devices that indicate an interest. This feature is supported on a spine-and-leaf topology where the network can consist of a mix of devices that support selective multicast Ethernet and those that do not support this feature.

  [See Selective Multicast Forwarding.]

- **BPDU protection in EVPN-VXLAN (QFX5100, QFX5110, and QFX5200 switches)**—Starting in Junos OS Release 18.4R2, you can enable BPDU protection in an EVPN-VXLAN configuration. With a spanning tree protocol configured on an edge port, you can enable BPDU protection. If a BPDU is received on the edge port, the edge port is disabled and it stops forwarding all traffic. You can also configure BPDU protection on VXLAN interfaces without a spanning tree protocol configured, or enable BPDU protection and have other traffic forwarded. Only the BPDUs are dropped, and all other traffic is forwarded. Additionally, you can unblock an interface either automatically or manually.

  - To enable BPDU protection with RSTP on an edge port on access and leaf devices:

    set protocols rstp interface *interface-name* edge

    set protocols rstp bpdu-block-on-edge

  - To enable BPDU protection with a spanning tree protocol on access and leaf devices:

    set protocols layer2-control bpdu-block interface *interface-name*

  - To enable BPDU protection but still forward other traffic on access and leaf devices:

    set protocols layer2-control bpdu-block interface *interface-name* drop

  - To automatically unblock an interface using an expiry timer on access and leaf devices:

    set protocols layer2-control bpdu-block disable-timeout *time in seconds*

  - To manually unblock an interface on access and leaf devices:

    run clear error bpdu interface all

- **Assisted replication in data centers with EVPN-VXLAN overlay networks (QFX Series switches)**—Starting in Junos OS Release 18.4R2, QFX Series switches support assisted replication (AR) in data centers with EVPN-VXLAN networks to optimize replication of BUM traffic being forwarded into the EVPN core.

Instead of flooding BUM traffic using ingress replication, devices configured as AR leaf devices forward the traffic to an AR replicator device that can better handle the replication load, and only the AR replicator device replicates and forwards the traffic to the overlay tunnels. You can configure switches in the QFX10000 line as AR replicator devices and any QFX Series devices that support EVPN-VXLAN as AR leaf devices.

AR devices advertise EVPN Type 3 (Inclusive Multicast Ethernet Tag [IMET]) routes that include special AR Type and Flags fields indicating AR device roles. The network can also include devices that do not support AR (regular network virtualization edge (RNVE) devices), which ignore AR routes and use ingress replication to forward BUM traffic toward the EVPN core.

You can configure AR with IGMP snooping to further optimize BUM traffic replication and forwarding.

To enable assisted replication and configure devices into AR replicator or AR leaf roles, use the **assisted-replication** configuration statement at the **[edit protocols evpn]** hierarchy level.

Software Defined Networking

- **Layer 2 and Layer 3 VXLAN gateways (EX4650 and QFX5120 switches)—**Starting with Junos OS Release 18.4R2, you can deploy EX4650 and QFX5120 switches as follows:

  - As a Layer 2 VXLAN gateway, or a Layer 2 and Layer 3 VXLAN gateway in an EVPN overlay network
  - (QFX5120 switches only) As a Layer 2 VXLAN gateway in an OVSDB overlay network

  VXLAN is an overlay technology that allows you to stretch Layer 2 connections over an intervening Layer 3 network by encapsulating (tunneling) Ethernet frames in a VXLAN packet that includes IP addresses. Using VXLANs to connect Layer 2 domains over a Layer 3 network means that you do not need to use the Spanning Tree Protocol (STP) to converge the topology (so no links are blocked) but can use more robust routing protocols in the Layer 3 network instead.

  [See Understanding VXLANs.]

- **OVSDB support with VMware NSX for vSphere (QFX5120 switches)—**Starting with Junos OS Release 18.4R2, the Open vSwitch Database (OVSDB) management protocol provides a control plane through which an NSX controller can provision QFX5120 switches. In an environment in which NSX Release 6.4.5 or later is deployed, an NSX controller and these switches can exchange control and statistical information, thereby enabling virtual machine (VM) traffic from entities in a virtualized network to be forwarded to entities in a physical network and the reverse.

  The physical underlay network over which OVSDB-VXLAN is commonly deployed is a two-layer IP fabric that includes spine and leaf devices. The spine devices function as Layer 3 VXLAN gateways, and the leaf devices function as Layer 2 VXLAN gateways. You can deploy QFX5120 switches as leaf devices in this fabric.

  [See Understanding the OVSDB Protocol Running on Juniper Networks Devices.]

New and Changed Features: 18.4R1

Authentication, Authorization, and Accounting (AAA)

- **Support for password change policy enhancement (QFX Series)**—Starting in Junos OS Release 18.4R1, the Junos OS password change policy for local user accounts is enhanced to comply with additional password policies. As part of the policy improvement, you can configure the following:
  - **maximum-lifetime-value**—The maximum duration of a password. The password expires after the maximum is reached.
  - **minimum-lifetime-value**—The minimum duration of a password. You cannot change the password until the minimum duration is reached.

  [See password.]

Class of Service (CoS)

- **Class of service support on VXLAN interfaces (QFX10000)**—Starting with Junos OS 18.4R1, standard class of service (CoS) features--classifiers, rewrite rules, and schedulers--are supported on VXLAN interfaces on the QFX10000 line of switches.

  [See Understanding CoS on OVSDB-Managed VXLAN Interfaces.]

- **Class of service support on VXLAN interfaces (QFX5100)**—Starting with Junos OS 18.4R1, standard class of service (CoS) features - classifiers, rewrite rules, and schedulers - are supported on VXLAN interfaces on QFX5100 switches.

  [See Understanding CoS on OVSDB-Managed VXLAN Interfaces.]

EVPNs

- **Support for graceful restart on EVPN-VXLAN (QFX Series)**—Starting in Junos OS Release 18.4R1, Junos OS supports graceful restart on EVPN-VXLAN on EX9200 and QFX Series switches and MX Series routers. Graceful restart allows the device to recover from a routing process restart or Routing Engine switchover without nonstop active routing (NSR) enabled.

  [See NSR and Unified ISSU Support for EVPN Overview.]

- **Selective multicast forwarding and SMET support in EVPN-VXLAN (QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 18.4R1, Junos OS supports selective multicast forwarding in a centrally EVPN-VXLAN network. Devices on a bridge domain with IGMP snooping enabled will monitor traffic on the access interfaces and selective forwarding towards the core. Devices that support selective multicast forwarding replicate and forward multicast traffic only to other interested devices. This feature is supported on a centrally-routed spine-and-leaf topology on QFX 10000 switches where the network can consist of a mix of SMET supported and non-SMET supported devices. This is achieved because the ingress devices can flood multicast traffic to the non-SMET capable devices while selectively forwarding the traffic among SMET capable devices. The ingress device can determine whether a device on the EVPN network is capable of supporting SMET by the presence or absences of the multicast flag community

in a EVPN type 3 route message and will forward the traffic accordingly. Thus, the data center fabric can be upgraded in phases without disrupting existing multicast operations.

[See Selective Multicast Forwarding .]

- **Support for VMTO for ingress traffic (QFX Series)**—Starting in Junos OS Release 18.4R1, you can configure a leaf or spine device that is configured as a Layer 3 gateway to support virtual machine traffic optimization (VMTO) for ingress traffic. VMTO eliminates the unnecessary ingress routing to default gateways when a virtual machine is moved from one data center to another.

  To enable VMTO, configure **remote-ip-host** routes at the **[edit routing-instances routing-*instance-name* protocols evpn]** hierarchy level. You can also filter out the unwanted routes by configuring an import policy under the **remote-ip-host routes** option.

  [See Configuring EVPN Routing Instances.]

- **Support for multihomed proxy advertisement (QFX Series)**—Starting in Junos OS Release 18.4R1, Junos OS now provides enhanced support to proxy advertise the MAC address and IP route entry from all leaf devices that are multihomed to a CE device. This can prevent traffic loss when one of the connection to the leaf device fails. To support the multihomed proxy advertisement, all multihomed PE devices should have the same multihomed proxy advertisement bit value. The multihomed proxy advertisement feature is enabled by default, and Junos OS uses the default multihomed proxy advertisement bit value of 0x20.

  [See EVPN Multihoming Overview.]

- **Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface (QFX5100, QFX5110, and QFX5200 switches)**—You can configure and commit the following on a physical interface of a QFX5100, QFX5110, or QFX5200 switch in an EVPN-VXLAN environment:

  - Layer 2 bridging (**family ethernet-switching**) on any logical interface unit number (unit 0 and any nonzero unit number).
  - VXLAN on any logical interface unit number (unit 0 and any nonzero unit number).
  - Layer 2 bridging (**family ethernet-switching** and **encapsulation vlan-bridge**) on different logical interfaces (unit 0 and any nonzero unit number).
  - Layer 3 IPv4 routing (**family inet**) and VXLAN on different logical interfaces (unit 0 and any nonzero unit number).

  For these configurations to be successfully committed and to work properly, you must specify the **encapsulation flexible-ethernet-services** configuration statement at the physical interface level—for example, **set interfaces xe-0/0/5 encapsulation flexible-ethernet-services**.

  This feature was previously introduced in Junos OS Release 18.1R3.

  [See Understanding Flexible Ethernet Services Support With EVPN-VXLAN.]

- **Automatically generated Ethernet segment identifiers in EVPN-VXLAN and EVPN-MPLS networks (MX240, MX480, QFX5100, and QFX5110)—**Starting in Junos OS Release 18.4R1, you can configure aggregated Ethernet interfaces and aggregated Ethernet logical interfaces to automatically derive Ethernet

segment identifiers (ESIs) from the Link Aggregation Control Protocol (LACP) configuration. This feature is supported in the following environments:

- On Juniper Networks devices that are multihomed in active-active mode in an EVPN-VXLAN overlay network.
- On Juniper Networks devices that are multihomed in active-standby or active-active mode in an EVPN-MPLS overlay network.

[See Understanding Automatically Generated and Assigned ESIs in EVPN Networks.]

- **MAC filtering, storm control, and port mirroring support in EVPN-VXLAN overlay networks (QFX5100 and QFX5110 switches)**—QFX5100 and QFX5110 switches support the following features in an EVPN-VXLAN overlay network:
  - MAC filtering
  - Storm control
  - Port mirroring and analyzers

  [See MAC Filtering, Storm Control, and Port Mirroring Support on EVPN-VXLAN Interfaces.]

- **MAC filtering and storm control support in EVPN-VXLAN overlay networks (QFX10002 and QFX10008 switches)**—QFX10002 and QFX10008 switches support the following features in an EVPN-VXLAN overlay network:
  - MAC filtering
  - Storm control

  [See MAC Filtering, Storm Control, and Port Mirroring Support on EVPN-VXLAN Interfaces.]

- **IPv6 data traffic support through an EVPN-VXLAN overlay network (QFX10000 and QFX5110 switches)**—Starting with Junos OS Release 18.4R1, QFX10000 and QFX5110 switches that function as Layer 3 VXLAN gateways can route IPv6 data traffic through an EVPN-VXLAN overlay network. With this feature enabled, Layer 2 or 3 data packets from one IPv6 host to another IPv6 host are encapsulated with an IPv4 outer header and transported over the IPv4 underlay network. The Layer 3 VXLAN gateways in the EVPN-VXLAN overlay network learn the IPv6 routes through the exchange of EVPN type-2 and type-5 routes.

  This feature was previously introduced in Junos OS Release 15.1X53-D30 on QFX10000 switches.

  [See Routing IPv6 Data Traffic through an EVPN-VXLAN Network With an IPv4 Underlay.]

High Availability (HA) and Resiliency

- **VRRP scale improvements per aggregated Ethernet bundle (QFX Series)**—Starting in Junos OS Release 18.4R1, you can configure up to 4000 active VRRP sessions per aggregated Ethernet bundle on QFX Series routers. To configure VRRP support, include the **vrrp-group** statement at the **[edit interfaces *interface-name* unit *logical-unit-number* family inet address *ip-address*]** hierarchy level.

  [See Understanding VRRP]

Junos on White Box

- **Junos on White Box**—Starting with Junos OS Release 18.4R1, the Junos on White Box software provides a disaggregated Junos that decouples the Junos operating system from Juniper Networks switches and runs as independent software on Open Compute Project (OCP)-compliant network hardware, enabling you to use that hardware in your data center (DC) networks and providing a robust, feature-rich network operating system for enabling the DC Fabric buildout. Junos for White Box is standalone software providing standards-based network protocols such as ISIS and BGP, overlay technology such as VXLAN with EVPN control plane, and full automation capabilities and is similar to the reliable, high performance Junos OS that powers the Juniper Networks QFX Series Data Center portfolio.

  Key Junos OS features that enhance the functionality and capabilities of the White Box switches include:

  - Software modularity, with process modules running independently in their own protected memory space and with the ability to do process restarts.
  - Uninterrupted routing and forwarding, with features such as nonstop active routing (NSR) and nonstop bridging (NSB).
  - Commit and rollback functionality that ensures error-free network configurations.
  - A powerful set of scripts for on-box problem detection, reporting, and resolution.

  > **NOTE:** The feature above was previously introduced in Junos OS Release 18.1R3.

  [See Junos on White Box Documentation.]

  The following features are supported in Junos on White Box in Junos OS Release 18.4R1:

  - Class of service (CoS) support. [See Overview of Junos OS CoS.]
  - Layer 2 VXLAN gateway and EVPN control plane and VXLAN data plane support. [See Understanding VXLANs; Understanding EVPN with VXLAN Data Plane Encapsulation.]
  - Multichassis link aggregation (MC-LAG). [See Understanding Multichassis Link Aggregation Groups.]
  - IPv4 GRE support. [See Understanding Generic Routing Encapsulation.]
  - Link aggregation and resilient hashing support. [See Understanding the Use of Resilient Hashing to Minimize Flow Remapping in Trunk/ECMP Groups.]
  - Channelizing Ethernet interfaces support. [See Channelizing Interfaces on Switches.]

- IPv6 protocols, including Neighbor Discovery Protocol; Virtual Router Redundancy Protocol (VRRP) for IPv6; Protocol Independent Multicast (PIM) for IPv6; BGP, IS-IS, and OSPFv3 for IPv6; unicast IPv6 for virtual-router instances; and DHCPv6. [See Example: Configuring IPv6 Interfaces and Enabling Neighbor Discovery; Verifying and Managing DHCPv6 Relay Configuration.]
- Layer 2 features: VLAN support; Link Layer Discovery Protocol (LLDP) support; Q-in-Q tunneling support; Spanning Tree Protocol (STP), Rapid Spanning Tree Protocol (RSTP), Multiple Spanning Tree Protocol (MSTP), and VLAN Spanning Tree Protocol (VSTP) support. [See Ethernet Switching User Guide.]
- Private VLANs (PVLANs)—including PVLANs with IRB interfaces—support. [See Understanding Private VLANs.]
- MPLS support. [See MPLS Overview.]
- Hierarchical ECMP and ECMP support on LSR. [See Overview of Hierarchical ECMP Groups; Configuring ECMP Next Hops for RSVP and LDP LSPs for Load Balancing.]
- Layer 2 and Layer 3 multicast support. [See Multicast Configuration Overview.]
- Junos Telemetry Interface (JTI) support. [See Overview of the Junos Telemetry Interface.]
- Services support: sFlow, analyzers/port mirroring, including remote port mirroring to an IP address (GRE encapsulation). [See Overview of sFlow Technology; Understanding Port Mirroring.]
- Firewall filter support and policers and counters support.

  [See Overview of Firewall Filters; Policer Implementation Overview.]
- Layer 3 unicast routing protocol support. [See BGP User Guide; IS-IS User Guide; OSPF User Guide; Protocol-Independent Routing Properties User Guide; RIP User Guide.]
- Access security features support. [See Overview of sFlow Technology; Understanding Port Mirroring.]
- Storm control support. [See Understanding Storm Control.]
- Distributed denial of service (DDoS) protection support. [See Distributed Denial-of-Service (DDoS) Protection Overview.]
- Open Network Install Environment (ONIE) support. [See Installing and Recovering Software Using the Open Network Install Environment (ONIE).]
- Zero Touch Provisioning (ZTP) support. [See Zero Touch Provisioning.]
- Support for Converged Enhanced Ethernet (CEE) features. [See Traffic Management User Guide for the QFX Series and EX4600 Switches.]

NOTE: The features above were previously introduced in Junos OS Release 18.1R3.

- **Layer 2 and 3 families, encapsulation types, and VXLAN on the same physical interface (Junos on White Box)**—You can configure and successfully commit the following on a physical interface of a switch in an EVPN-VXLAN environment:
  - Layer 2 bridging (**family ethernet-switching**) on any logical interface unit number (unit 0 and any nonzero unit number).
  - VXLAN on any logical interface unit number (unit 0 and any nonzero unit number).
  - Layer 2 bridging (**family ethernet-switching** and **encapsulation vlan-bridge**) on different logical interfaces (unit 0 and any nonzero unit number).
  - Layer 3 IPv4 routing (**family inet**) and VXLAN on different logical interfaces (unit 0 and any nonzero unit number).

  For the above configurations to be successfully committed and work properly, you must specify the **encapsulation flexible-ethernet-services** configuration statements at the physical interface level—for example, **set interfaces xe-0 /0/5 encapsulation flexible-ethernet-services**.

  This feature was previously introduced in Junos OS Release 18.1R3.

  [See Understanding Flexible Ethernet Services Support With EVPN-VXLAN.]

- **Automatically generated Ethernet segment identifiers in EVPN-VXLAN networks (Junos on White Box)—**Starting in Junos OS Release 18.4R1, you can configure aggregated Ethernet interfaces and aggregated Ethernet logical interfaces to automatically derive Ethernet segment identifiers (ESIs) from the Link Aggregation Control Protocol (LACP) configuration. We support this feature on switches that are multihomed in active-active mode in an EVPN-VXLAN network.

  [See Understanding Automatically Generated and Assigned ESIs in EVPN Networks.]

Operation, Administration, and Maintenance (OAM)

- **Connectivity fault management (CFM) support (QFX5200 and QFX5210)**—IEEE 802.1ag CFM provides fault isolation and detection over large Layer 2 networks that may span several service provider networks. You can configure CFM to monitor, isolate, and verify faults in these interconnected provider bridge networks. Starting in Junos OS Release 18.4R1, Junos OS provides CFM support on QFX5200 and QFX5210.

  CFM support on QFX5200 and QFX5210 has the following limitations:
  - CFM support is provided via software using filters. This can impact scaling.
  - Inline Packet Forwarding Engine mode is not supported. In Inline PFE mode, you can delegate periodic packet management (PPM) processing to the Packet Forwarding Engine which results in faster packet handling. The CCM interval supported is 10 milliseconds.
  - Performance monitoring (ITU-T Y.1731 Ethernet Service OAM) is not supported.
  - CCM interval of less than 1 second is not supported.
  - CFM is not supported on routed interfaces and aggregated Ethernet (lag) interfaces.

- MIP half function, to divide the MIP functionality into two unidirectional segments to improve network coverage, is not supported.
- Up MEP is not supported.
- Total number of CFM sessions supported is 20.

[See Understanding Ethernet OAM Connectivity Fault Management for Switches.]

System Management

- **Passive Monitoring support (QFX10000 switches)**— Starting with Junos OS Release 18.4R1, you can enable passive monitoring on the switch to passively capture traffic from monitoring interfaces. Passive monitoring provides filtering capabilities for monitoring ingress and egress traffic at the Internet point of presence (PoP) where security networks are attached. With passive monitoring, the switch does not route packets from the monitored interface or run any routing protocols related to those interfaces. It only receives traffic flows, collects intercepted traffic, and exports it to monitoring tools like IDS servers and packet analyzers, or other devices such as routers or end node hosts. To enable this feature, include the **passive-monitor-mode** statement at the **[edit interface]** hierarchy level. This feature was previously supported in an "X" release of Junos OS.

  See [Understanding Passive Monitoring on QFX10000 Switches.]

- **IPv6 support added to Precision Time Protocol (PTP) G.8275.2) enhanced profile (QFX5110 and QFX5200 switches)**— Starting with Junos OS Release 18.4R1, the G.8275.2 enhanced profile supports IPv6 transport.

  To configure the G.8275.2 enhanced profile, enable the **g.8275.2.enh** statement at the **[edit protocols ptp profile-type]** Junos OS CLI hierarchy.

  To configure IPv6 transport, enable the **ipv6** statement at the **[edit protocols ptp master interface interface-name unicast-mode transport]** and **[edit protocols ptp slave interface interface-name unicast-mode transport]** Junos OS CLI hierarchies.

VPNs

- **Support to control traceroute over Layer 3 VPN (QFX Series)**—Starting in Junos OS Release 18.4R1, in a Layer 3 VPN topology with **vrf-table-label** configured and multiple customer edge (CE) routers configured in the same VPN routing and forwarding (VRF) routing instance, when traceroute is performed to a remote provider edge (PE) router for a CE-facing network, the ICMP time exceeded packet determines the correct IP address as the source address.

  To control the traceroute over Layer 3 VPN topology with vrf-table-label configured and multiple CE routers configured in the same VRF, you can configure **allow-l3vpn-traceroute-src-select** at the**[edit system]** hierarchy level that determines the correct IP source address by reviewing the destination routing instance and destination IP address.

  [See allow-l3vpn-traceroute-src-select.]

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 18.4R2 for the QFX Series.

Release 18.4R2-S3 Changes in Behavior and Syntax

Platform and Infrastructure

- **Logical Interface is created along with physical Interface by default (EX Series switches, QFX Series switches, MX Series routers)**—The logical interface is created on ge, et, xe interfaces along with the physical interface, by default. In earlier Junos OS Releases, by default, only physical interfaces were created. For example, for ge interfaces, earlier when you view the **show interfaces** command, by default, only the physical interface (ge-0/0/0), was displayed. Now, the logical interface (ge-0/0/0.16386) is also displayed.

Changes in Behavior and Syntax: 18.4R2-S1

Software Defined Networking (SDN)

- **Increase in the maximum value of delegation-cleanup-timeout (QFX Series)**—You can now configure a maximum of *2147483647* seconds as the delegation cleanup time for a Path Computation Client (PCC). This extends the time taken by the PCC to retain the last provided path over a PCEP session from the last session down time.

  With the increase in maximum value of **delegation-cleanup-timeout** from *600* to *2147483647* seconds, you can benefit during a Path Computation Element (PCE) failover, or other network issues that may disrupt the PCEP session with the main active stateful PCE.

  [See delegation-cleanup-timeout.]

Changes in Behavior and Syntax: 18.4R2

IN THIS SECTION

EVPNs

- **New options in show evpn instance command (QFX series)**—Starting in Junos OS Release 18.4R2, 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.

- **Changes to show evpn instance extensive command (QFX series)**—Starting in Junos OS Release 18.4R2, the output for **show evpn instance extensive** displays information on the core next hop for unknown multicast streams only. For known multicast streams, use the **show evpn igmp-snooping proxy** command.

Interfaces and Chassis

- **Commit error when GRE interface and tunnel source interface configured in different routing instances (QFX Series)**—In Junos OS Releases 17.3R4, 17.4R3, 18.1R4, 18.2R3, 18.3R2, and 18.4R2, QFX Series switches do not support configuring a 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 .]

- **New XML tag element <lacp-hold-up-state> added in show lacp interfaces XML display (QFX Series)**—In Junos OS Release 18.4R2, 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 state from 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. This information for each interface is now 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 18.4R2, the **resilient-hash** statement is no longer available at the **[edit interfaces ae*x* aggregated-ether-options**] hierarchy level. Resilient hashing is not supported on LAGs on QFX5200 and QFX5210.

  [See aggregated-ether-options.]

- **Logical interface is created along with physical interface by default (QFX10000 and QFX5000 line of switches)**—In Junos OS Release 18.4R2, on the QFX10000 line of switches, by default, logical interface are created on et-, sxe-, and non-channelized xe- interface along with the physical interface. In earlier Junos OS Releases, by default, only physical interfaces are created.

  On QFX5000 line of routers, by default logical interface is created on channelized xe- interfaces. In earlier Junos OS releases, by default, channelized interfaces (xe-0/0/0:1, xe-0/0/0:2, and so on) do not have logical interfaces by default and only the nonchannelized et- and xe- interfaces and sxe- creates logical interfaces.

Security

- **Syslog or log action on firewall drops packets (QFX5000 switches)**—Starting in Junos OS Release 18.4R2, 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 18.4R2, a warning message is displayed whenever a firewall term includes **log** or **syslog** with the **accept** filter action.

Changes in Behavior and Syntax: 18.4R1

IN THIS SECTION

Interfaces and Chassis

- **Change in default action for fatal errors (QFX10002, QFX10008, and QFX10016 switches)**—Starting in Junos OS Release 18.4R1, by default, for all fatal errors on the QFX10000 line of switches, Junos OS raises an alarm and disables all Packet Forwarding Engine interfaces that raised the error.
- **Support for creating layer 2 logical interface independently (QFX Series)**—In Junos OS Releases 18.4R1, 18.4R2, and later, QFX Series switches support creating layer 2 logical interface independent of layer 2 routing instance type. That is, you can configure and commit the layer 2 logical interfaces separately and add the interface to bridge-domain or Ethernet VPN (EVPN) routing instance separately. Note that the layer 2 logical interfaces works fine only when the interface is added to bridge domain or EVPN routing instance.

  In the earlier Junos OS releases, when an layer 2 logical interface configuration (units with encapsulation vlan-bridge configuration) is used, then the logical interface must be added as part of a bridge-domain or EVPN routing instance for the commit to succeed.

Network Management and Monitoring

- **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 18.4R1, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, the server must not return an RPC reply that encloses both an **<rpc-error>** element and an **<ok/>** element. If the operation is successful, but the server reply would enclose one or more **<rpc-error>** elements of severity warning in addition to the **<ok/>** element, then 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 encloses both an **<rpc-error>** element of severity warning and an **<ok/>** element.

SEE ALSO

Known Behavior

IN THIS SECTION

This section lists known behavior, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 for the QFX Series.

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

Class of Service (CoS)

- On QFX5120 switches, if the CoS configurations are modified when egress traffic shaped at very low rate (less than 50 Mbps), packets might get stuck in the MMU buffers permanently. It might cause ingress or egress traffic drops. When low rate shapers (less than 50 Mbps) are applied on egress queues, it is suggested to deactivate shaping before any CoS modification or ensure traffic is stopped before doing CoS modification. PR1367432

EVPN

- When a VLAN uses an IRB interface as the routing interface, the **vlan-id** parameter must be set to **none** to ensure proper traffic routing. This issue is platform independent. PR1287557

General Routing

- Port LEDs on QFX5100 do not work. If a device connects to a port on QFX5100, the port LED stays unlit. PR1317750
- Based on the memory availability, the QFX10002 can scale up to 300 remote PE devices with a total of 600 tunnels. To avoid exceeding memory,we recommend that you do not go beyond this scale. PR1329243
- When the sFlow collector can be reached only through the Routing Engine, because of heavy traffic, large samples can cause the Routing Engine CPU to become busy. PR1332337
- In an IP-CLOS topology, when a spine device and leaf device is rebooted, you might see a traffic loss for around 100 seconds. The reason for this is that, Junos OS starts advertising routes before Packet Forwarding Engine route programming is completed, which might cause traffic loss. PR1341398
- Hardware watchdog does not work on QFX10008 and QFX10002-60C/PTX10002-60C. PR1343131
- When a VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter will not be installed. PR1362609
- A few error messages related to function rt_mesh_group_add_check() will be seen during reboot and are harmless. PR1365049
- Autochannelization is not supported for 40GBASE-BXSR, QSFP+40GE-LX4, QSFP-100G-PSM4, and 100GBASE-BXSR optics. PR1366103
- 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 message is seen when channelization is detected in Junos OS Release 18.1R3. PR1366137
- When the egress-to-ingress option is enabled to use ingress TCAM for the egress filters, it is expected that the egress counters will count the packets on the ingress side as well. PR1369048

- Error logs are expected when routes pointing to the target next hop, which in turn points to the HOLD next hop. These error logs are present for short time. Later, when the next hop changes from HOLD next hop to valid next hop, unilist next hops will be walked again and updated with the appropriate weight and reroute counters. and no more error logs will be seen. PR1387559
- On Junos OS Release 18.4R1, an intermittent traffic loss is observed with RTG streams while flapping the RTG primary interface. PR1388082
- Re-ARP request sent without VLAN-ID (so Routing Engine ARP fails). PR1390794
- On QFX5000 devices, there is a possibility for the system to go to DB prompt in a reboot scenario. This is because of a known issue in the QEMU version in WRL7. PR1411826
- If the commit fails with **statements constraint check failed** even though the dependant configuration is in place, it is possible that main and dependent configurations are configured through different groups. It is due to system constraints. PR1437047

Layer 2 Features

- Currently, the **show multicast snooping route extensive** command is not supported on QFX Series devices. PR1386905
- In MH scenarios, QFX5 does not support transition of the remote learnt MAC (DR) to locally learnt MAC (DL) when the traffic hashes to MH PE where the MAC is programmed as DR. Because of this during MAC or MAC-IP aging, the MAC entry on both the PE devices will be deleted and re-learnt. PR1419988

MPLS

- There will not be any warning message about a Packet Forwarding Engine restart when MPLS tunnel extend configuration is deleted. PR1394722

Routing Protocols

- On QFX5120 platforms, 254 neighbors and 200,000 routes can be scaled for IS-ISv4. Beyond 200,000 routes with 254 neighbour, adjacency flaps and traffic drop will be seen. PR1368106
- Targeted broadcast functionality with VXLAN is not supported on QFX5000 platforms. For a non-VXLAN case, broadcast destination IP look up results in next hop with destination MAC of all 0xffs and gives the class-id for IFP to match and action to redirect to IPMC with VLAN membership check. In VXLAN case, l3 egress intf, egr l3 next hop, ingress l3 entry creations fail. PR1397086
- QFX10002: After applying firewall family ethernet-switching filter from ether-type arp, the firewall did not filter the ARP request and counter does not increment. The configuration works if we disable user-clan-id match from the term. PR1426590

User Interface and Configuration

- **Auto-complete caution for QFX10002-60C and PTX10002-60C personalities**—Starting in Junos OS Release 18.4R1, for QFX10002-60C and PTX10002-60C personalities, do not use auto-complete to display the list of arguments for the **request system software delete** command. You must look for the package name using the **show system software** command and then explicitly type the software package name in the **request system software delete** command.

  [See request system software delete].

Virtual Chassis

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

SEE ALSO

New and Changed Features | 222

Changes in Behavior and Syntax | 233

Known Issues | 240

Resolved Issues | 249

Documentation Updates | 262

Migration, Upgrade, and Downgrade Instructions | 263

Product Compatibility | 277

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software for the QFX Series switches in Junos OS Release 18.4R2.

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

EVPN

- The **mac-move-shutdown** statement stops working if a physical loop is introduced continuously in a quick succession of 10 minutes. The issue is not seen every time but might occur only if the physical loop is introduced at least four times. If the loops span a longer period, the issue is not seen. A test is performed to check the overall impact on basic features. There is no issue seen on basic learning or major impact on any protocol. This is a negative scenario, but it is unlikely to occur in a customer network where the multiple loops occur within a short time span. PR1284315
- On EVPN-VXlAN, ping over Type 5 IRB fails if multiple IRB are configured in routing instance along with loopback lo.0 interface also configured. PR1346894
- At times, when l2ald is restarted, a race condition occurs where a VTEP notification comes in from the kernel before lo0. As a result, l2ald is unable to process the VTEP add request and gets stuck in an indefinite loop. PR1384022
- On a QFX10,000 with **nonstop-routing** enabled and EVPN running, if Routing Engine switchover occurs, the EVPN traffic might have significant traffic loss. PR1394099
- [evpn_vxlan] [virtual_switch] IRB MAC/IP information is deleted from the Ethernet-switching ARP/ND table when **no-arp-suppression** is configured. PR1394959
- To filter and see the output of a specific ESI or neighbor information of an EVPN instance, we created two new choices, namely **show evpn instance <> esi-info esi <>** and **show evpn instance <> neighbor-info neighbor <>**. PR1402175
- In an assisted replication (AR) enabled network, the multicast traffic might silently get dropped and discarded toward AR-Leaf devices that do not support snooping if the AR replicators are snooping enabled. PR1403292
- ARP and IPv6 neighbor entries cannot be cleared when they are learned from EVPN multihome ESI. The following commands will not clear ARP and IPv6 neighbor entries when they are learned from EVPN multi-home ESI:

- **clear ethernet-switching evpn arp-table**
- **clear ethernet-switching evpn nd-table**
- **clear ethernet-switching mac-ip-table**. PR1446957

General Routing

- L3 multicast traffic does not converge 100 percent and continuous drop in traffic is observed after bringing down or up the downstream interface or while an FPC comes online after FPC restart. This happens with multicast replication for 1000 VLANs or IRB interfaces. PR1161485
- When you use **request system reboot**, the device undergoes zeroize, which triggers ZTP. During the mounting stage, **/var/db/scripts/import** does not get created, which later causes the configuration to be committed partially. This is seen in the warning **Warning: Commit failed, activating partial configuration. Warning: Edit the router configuration to fix these errors**. PR1289782
- Every load override and rollback operation increases the refcount by 1 and after it reaches the maximum value (65,535), the mgd process is terminated. When mgd is terminated, the active lock might remain, preventing any further commits. PR1313158
- Port LEDs on QFX5100 do not work. If a device connects to a port on QFX5100, the port LED stays unlit. PR1317750
- Interface uptime has increased by 8 seconds from Junos OS Release 17.4R1 to Junos OS Release 18.1R1. Also, SDK upgrade across releases might impact the parameters such as login prompt appear time, FPC up time, and interface up time after switch reboot. PR1324374
- On the QFX10002-60C, the filter operation with the log action is not supported for protocols other than Layer 2, IPv4, and IPv6. The following message is seen in the firewall logs: **Protocol 0 not recognized**. PR1325437
- BFD session over aggregated Ethernet flaps when a member link carrying the BFD Tx flaps. PR1333307
- On QFX10002, QFX10008, and QFX10016, ND is incorrectly working on IRB/Layer 3 interface with a discard filter. PR1338067
- On QFX10000 platforms, NETCONF over SSH traffic through TCP port 830 might hit the host path queue that is unclassified. This can result in DDoS violations in the unclassified queue. PR1345744
- QFX10000 platform drops the Aruba wireless Access Point (AP) heartbeat packets; as a result the Aruba wireless AP cannot work. PR1352805
- When MC-LAG is configured with **force-up** enabled on MC-LAG nodes, the LACP admin key should not match the key of the access or CE device. PR1362346
- On QFX5000 platforms, if lcmd is restarted, a chassisd core file is generated with traffic drop for few seconds. PR1363652

- On the QFX5100, if a scaled configuration involving a LAG interface, more that 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, FXPC might consume high CPU resources. No other system impact is observed. PR1363896
- On QFX5210, a filter with the routing instance applied to a family **inet** logical interfaces causes traffic to be discarded on unrelated interfaces PR1364020
- From Junos OS Release 17.3R1 and later, on the QFX10002 platform, in a rare condition, the IPFIX flow statistics (packet/byte counters) are incorrect in the exported record. Because the statistics are not collected properly, the flow might timeout and get deleted because of the inactive time out, 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 a 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 message is 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 is seen: **expr_nh_fwd_get_egress_install_mask:nh type Indirect of nh_id: # is invalid**. PR1367121
- A dedicated minimum number of buffers are reserved for some queues according to the Junos OS working model. These buffers are always available to those queues irrespective of the traffic pattern throughout the system. When the **clearing stat** statement is used, these values are visible. This cosmetic or minor issue has no functional impact. PR1367978
- If both the local and remote ends are auto-channelized and the local port QSFP transceiver is removed, then the 100-Gigabit Ethernet interface does not come up on port 62 after removing the SFP transceiver on port 30, which is channelized. PR1370887
- The DSCP values for IPv6 PTP packets exiting a QFX5110 are set as 111000 and go out only in the network control queue. PR1371064
- Changing the bridge-domain name breaks the communication for that particular bridge domain. PR1371495
- The L2 bridge domain might fail to be created on the Packet Forwarding Engine after the VLAN configuration is changed. For example, suppose there are three VLANs V1001, V1002, and V1003. V1001 is deleted and V1002's VLAN-ID and VNI is changed to that of V1001 and a new vlan V1200 is added with the VLAN-ID and VNI of VLAN V1002. After these changes, V1200 is not created in the Packet Forwarding Engine and the other two VLANs are functioning as expected. The reason why the new VLAN is not created is out of order messages. This is a timing issue. PR1371611

- MAC learning does not happen after restart of l2-learning daemon for interfaces on backup. Traffic still gets forwarded. PR1372220
- Static speed 100-Mbps setting remains after changing speed 100-Mbps to auto-negotiation. PR1372647
- USB upgrade of NOS image is not supported. PR1373900
- When CoS-based forwarding (CBF) is enabled, because of the indexed next-hop installation issue in the kernel, the rpd process might crash upon route flap and LSP flap. PR1374558
- In Junos OS Release 18.1R3, when one 50-Gigabit Ethernet port is taken down using the **ifconfig** command, the other one also goes down. PR1376389
- In a certain scenario where flows are sampled through aggregate bundles when J-Flow 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
- 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
- When the **show** command is taking a long time to display results, the STP might change states as BPDUs are no longer processed and cause multiple 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
- When 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
- Layer 3 gateway is not supported on QFX5110 with the SP style of configuration in Junos OS Release 18.1R3-S2 and Junos OS Release 18.4R1. PR1399131
- On QFX5100, traffic initiated from a server connected to an interface is 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
- In a system after deleting the scale L3 VXLAN configurations (4000 VLANs/VXLAN), LACP states for few interfaces might go to detached state. As a workaround, reconfigure affected aggregated interfaces. PR1406691
- PXE installation might fail in this release because of a failure in image upgrade post PXE initialization. PR1406743

- On a QFX5120 platform with the QSFP-100G-PSM4 transceiver, there is a possibility that because of the timing fault on field-programmable gate array (FPGA) hardware, the link might go down because the TX laser is disabled. PR1410687
- L2 logical interface configuration can now be committed separately from the bridge-domain or EVPN configuration. PR1414363
- On QFX5110 and QFX5120 platforms, uRPF check in strict mode does not work properly. PR1417546
- ERSPAN traffic is not tagged when the output interface is a trunk port. PR1418162
- libvirtMib_suba core files might be generated during the installation of images. There is no functional impact because of these core files, because the core file is generated in the libvirtMib_subagent. PR1419536
- For transit static LSPs, QFX5120-48Y (trident 3 based)/QFX5120-32C, using broadcom trident3 devices might end up in swapping with an invalid label instead of POP/PHP action and might result in packet drop in the adjacent LER node. Because the TD3 chipset has additional capabilities for MPLS. This issue is applicable only to QFX5120-48Y (trident 3 based)/QFX5120-32C, using broadcom trident3 platforms and not applicable to other platforms. As a workaround, removing/re-applying the static transit LSP configuration. PR1420370
- When using the QSFP-100G-PSM4 transceiver on a QFX5120, there is a possibility that after leaving the setup idle more than 10 days, the port might not be available as it goes to channelized state and gets stuck there. PR1424647
- QFX5100-96F and QFX5100-VC platforms might be unable to commit baseline configuration after zeroization. **{master:0}[edit] root# commit check Mar 26 05:50:48 mustd: UI_FILE_OPERATION_FAILED: File /var/run/db/enable-process.data doesn't exist Mar 26 05:50:48 mgd[1938]: UI_FILE_OPERATION_FAILED: Failed to open /var/run/db/enable-process.data+ file error: Failed to open /var/run/db/enable-process.data+ file error: configuration check-out failed: daemon file propagation failed**. PR1426341
- CRC errors might be seen when other manufacturer device is connected to QFX10000 on a 100G link with QSFP-100GBASE-LR4-T2. Other manufacturer device report CRC errors and input errors on those 100G links. The QFX10000 interfaces do not show any errors. It might cause packet loss. PR1427093
- More number of MAC or MAC-IP entries can be learnt if MAC or MAC-IP limit is configured in a particular sequence. An example is shown below:

  1. Learn 50 remote entries

  2. Configure MAC limit of 20 (remote entries remain intact, this works as expected)

  3. Learn 50 local entries. At this point, no local entries must be learnt. The MAC limit is 20.

  However, all 50 local MAC entries get learned causing the MAC count to be 100, which is incorrect. The same issue will be seen for MAC-IP. PR1428572
- Protocols get forwarded when using non-existing SSM map source address in IGMPv3 instead of pruning. PR1435648

- Because of PE chip limitation when the underlay is tagged after decapsulation when inner packet is recirculated it still retains the VLAN tag property from outer header since outer header is tagged. Thus, four bytes of inner tag is overwritten in inner packet and packet get corrupted. As a result, trap seen in EGP chksum in PE chip. As a workaround, enable **encapsulate-inner-vlan** configuration statement. PR1435864
- When LACP is configured with link protection and force-up on local device, and the peer is configured with link protection, disabling the active member on the peer device causes LACP MUX state to be stuck in attached state. PR1439268
- NDI cannot be used in VLAN with IRB on QFX5200. Neighborhood advertisements/ solicit packets destined to host are getting dropped with NDI inspection (under DHCPv6 security) on a VLAN with IRB configuration on QFX5200 in Junos OS Release 18.4 and later. PR1439844
- In a large-scale setup with 4000 VLANs and many VTEPs and access interfaces for each VLAN, when AR-leaf configuration is enabled or disabled, the fxpc process consumes high CPU for a long time. It might take around 30 minutes for all MAC and MAC+IP entries to be installed in the system. PR1442390
- Sometimes when the lo0 IP for a source VTEP is changed, there can be a stale entry of IP in the vlan_xlate table. So, if the traffic needs to be l3 routed to the another box having that IP, traffic might drop. PR1443390
- When there is only one term containing user-vlan-id match condition and there are no other terms in the IPACL_VXLAN filter except discard, the discard action for non-matching traffic will work for only that VLAN which is specified under user-vlan-id and not for other VXLAN VLANs which are part of that trunk port on which filter is applied. This can be ignored by adding another term to the filter which does not contain user-vlan-id match. PR1446489
- When the full configuration and the base configuration are loaded alternatively, sometimes the l3_intf created in the hardware might remain stale. When an arp_ndp entry is still using the interface, the interface might get deleted leading to a traffic loss. PR1441915

Infrastructure

- When there is a high route churn or when there is a high rate of route updates being pushed to the kernel, the **show interface** command might show a delay or not show all statistics due to route updates being prioritized over statistic messages. PR1250328

Interfaces and Chassis

- Traffic drop observed when trying to configure aggregated Ethernet interface description. PR1305794

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 QFX5000 Series platforms with spine-leaf scenario, when two or more than two underlay interfaces with ECMP are brought down on leaf devices, the multihop BFD overlay sessions between spines and leafs might flap. And if BFD flaps, the protocols depending on BFD (typically, IBGP protocols) might also flap, which leads to traffic impact. PR1416941

Layer 2 Features

- In QFX5000 platforms, when a scaled configuration (with more than 3000 bridge domains and more than 8000 ESI logical interfaces) is overwritten with a functional configuration (with 4 bridge domains and less than 10 ESI logical interfaces), using the **load override** command, it takes around 2 minutes for cleanup and adding of the new configuration. Without waiting for 2 minutes, the configuration is overwritten multiple times, then some bridge domains are not cleaned up in the CLI. PR1363410
- On a QFX5100 Q-in-Q interfaces might stop working for a certain **vlan-id-list** configured under a physical interface. This is the result of a Packet Forwarding Engine binary issue, which 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 the kernel. The kernel tries to remove the implicit filters, which return ENOENT, as the entries are already disabled during the first commit. The following messages are harmless to the system. PR1400606
- On QFX5000 platforms, the fxpc might keep crashing when a firewall filter is applied on a logical unit of a dsc-interface. This has traffic impact. PR1428350
- In a scenario where aggregated Ethernet has an IPACL_VXLAN filter applied and if all the child members of the same aggregated Ethernet are disabled and enabled back, only two hardware instances in the Packet Forwarding Engine can be seen. Therefore, double the packet count of incoming traffic is observed. This issue is seen when all the child members are deactivated and activated back. PR1441424

MPLS

- There could be some lingering RSVP state, which might 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 the 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, 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

Routing Protocols

- In an MC-LAG setup, when status-control standby is rebooting and status-control active is down, and if the ICCP session-establishment timer is configured less than or equal to the **init-delay-timer** on status-control standby, the mc-ae status of status-control standby might not become active until the peer node is up. To avoid this, during these cases, ICCP session-establishment timer should be configured greater than **init-delay-timer** with preferably 100s or more. PR1348648
- On QFX Series platforms, in a corner scenario with a Virtual Chassis setup, if storm control configuration is enabled on interfaces and multicast traffic ingresses on the interfaces, some storm control error logs might be observed on these interfaces. It is only seen in one customer setup and not reproducible in a local setup. Also, it is just a logging issue and has no traffic impact. PR1355607
- 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 QFX Series switches except for QFX10000 switches, if a host-destined packets (that is, the destination address belongs to the device) come from the interface with ingress filter of log/syslog action (that is, **filter <> term <> then log/syslog**), such packets should not be dropped and reach the Routing Engine. PR1379718
- On QFX5100 VC/VCF, the following error is observed: **BRCM_NH-,brcm_nh_bdvlan_ucast_uninstall(),128:l3 nh 6594 unintsall failed in h/w with Mini-PDT base configurations**. There is no functionality impact because of this error. message. PR1407175
- The separate group creation for egress-to-ingress feature (in QFX5110) will be supported from Junos OS Release 19.1R2 and later. In Junos OS Release 19.1R1, this feature will use the already existing

ERACL firewall group. As a result of this extra qualifier in the ERACL group, the group will operate in double wide mode instead of single wide, thus leading to reduced scale. PR1408670

- On QFX5110 and QFX5200 platforms, the dcpfe might crash if any interface flaps. PR1415297
- In BGP graceful restart scenario, including helper mode that is enabled by default, rdp generates a core file because of the improper handling of BGP graceful restart stale routes during the deletion of the BGP neighbor. The rpd crashes resulting in an impact on service and traffic. PR1427987
- In BGP graceful restart scenario, including helper mode that is enabled by default, the rpd process generates a core file because of the improper handling of BGP graceful restart stale routes during the deletion of the BGP neighbor. The rpd process crashes resulting in an impact on service and traffic. PR1441554

SEE ALSO

New and Changed Features | **222**

Changes in Behavior and Syntax | **233**

Known Behavior | **237**

Resolved Issues | **249**

Documentation Updates | **262**

Migration, Upgrade, and Downgrade Instructions | **263**

Product Compatibility | **277**

Resolved Issues

IN THIS SECTION

This section lists the issues fixed for the QFX Series switches in Junos OS Release 18.4R2 for QFX Series.

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

Resolved Issues:18.4R2

Class of Service (CoS)

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

EVPN

- The rpd process might crash with EVPN type 3 route churn. PR1394803
- 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
- ARP refresh functionality might fail in an EVPN scenario. PR1399873
- EVPN: In the non-collapsed (centralized) topology, when one of the two spines deactivates the underlay protocol (OSPF), the leaf still points the virtual gateway MAC next hop to the spine that is down. PR1403524
- ARP entry is still pointing to the failed VTEP after the PE-CE link fails for a 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
- Extra incorrect MAC move might be seen when the host moves continuously between the different ESIs PR1429821

Forwarding and Sampling

- On Junos OS, firewall filter terms named "internal-1" and "internal-2" are ignored. PR1394922

General Routing

- The 1iGigabit copper module interface shows **Link-mode: Half-duplex** on QFX10000 line platforms. PR1286709
- On QFX5120, convergence delay between PE1 and P router links is more than the expected delay value. PR1364244
- RIPv2 update packets might not be sent with **IGMP snooping** enabled. PR1375332
- EM policy update is needed on QFX5210-64C. PR1380077
- The overlay ECMP might not work as expected on QFX5110 in an EVPN-VXLAN environment PR1380084
- There is an inconsistency in applying a scheduler map with excess rate on the physical interface and aggregated Ethernet interface. PR1380294
- Traffic is silently dropped and discarded when the FPC is taken offline in an 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 restart. PR1383426
- DMA failure errors might be seen when the cache is flushed or the cache is full. PR1383608
- DHCP packets might be dropped in a Junos fusion data center scenario (QFX10000 line of devices). PR1383623
- Last reboot reason is not correct if the device is rebooted because of power cycling. PR1383693
- The Virtual Chassis could not come up after upgrading to QFX5E platforms (TVP-based platforms for QFX5100 or QFX5200 switches). PR1383876
- Disable reporting of correctable single-bit error on Hybrid Memory Cube (HMC) and prevent a major alarm. PR1384435
- QFX5120 occasionally two of the channelized 25-Gigabit ports using 4x25-Gigabit breakout cable will not come up after Junos OS reboot. PR1384898
- The spine EVPN routes might be stuck in a hidden state with the next hop as unusable after the FPC is offline in the spine. PR1386147
- The **show chassis errors active detail** command is not supported on QFK5000 platforms. PR1386255
- The rpd process might end up with stuck krt queue entries in a VRF scenario. PR1386475
- Traffic drop might be seen on QFX10000 platforms with EVPN-VXLAN configured. PR1387593
- QFX5100, QFX5110, QFX5200, and QFX5210 Virtual Chassis could not be formed normally. PR1387730
- On QFX5100 Virtual Chassis, ARP received on SP-Style interface is not sent to all RVTEPs. Normal BUM traffic works fine. PR1388811
- FPC might crash on QFX5100 platforms in a large-scale scenario PR1389872
- Input rate pps does not increase on QFX5200-48Y uplink ports when the packet is a pure L2 packet like non-etherII or non-EtherSnap. PR1389908
- An incorrect error message might be seen when Jflow sensors are configured with reporting rate less than 30 seconds. PR1390740
- 10-Gigabit copper link flapping might happen during a TISSU operation of QFX5100-48T switches. PR1393628
- IPv6 next hop programming issue might be observed on QFX10000 devices. PR1393937
- On QFX5110 Virtual Chassis, fan tray output is not displayed for backup Routing Engine. PR1394655
- PTP-over-Ethernet traffic could be dropped if IGMP and PTP TC are configured together. PR1395186
- Unable to install licenses automatically on QFX Series platforms. PR1395534

- **BRCM_NH-,brcm_bcm_mpls_tunnel_initiator_clear(),226:bcm_mpls_tunnel_initiator_get failed intf = 4 failure** error logs might seen in syslog. PR1396014
- The subscriber bindings might not be successful on QFX Series platforms. PR1396470
- On QFX5110, the fan LED turns amber randomly. PR1398349
- High jsd or na-grpcd CPU usage might be seen even when JET or JTI is not used. PR1398398
- CPU interrupt process is high because of the intr{swi4: clock (0)} on QFX5100-48T-6Q running a QFX 5e Series image and Junos OS 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 QFX10000 Series platform. PR1399369
- The DHCPv6 relay packets might be dropped by the DHCP relay. PR1399683
- SFP-LX10 does not work on QFX5110 PR1399878
- PEM I2C failure alarm might be shown incorrectly as failed. PR1400380
- MAC limit with persistent MAC is not working after reboot PR1400507
- Only one Packet Forwarding Engine might be disabled on an FPC with multiple PFEs in error/wedge condition. PR1400716
- The authd might crash when issuing **show network-access requests pending** command during the authd restart. PR1401249
- File permissions are changed for **/var/db/scripts** files after reboot. PR1402852
- The STP does not work when aggregated interfaces number is "ae1000" or above in QFX5000 and "ae480" or above in other QFX Series platforms. 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 will not be resolved if a native VLAN ID is configured for an LAG access interface. PR1404895
- Commit warning message occurs on QFX5100. PR1405138
- Executing the command **request system configuration rescue save** might fail with error messages. PR1405189
- DHCP does not work for some clients in Junos fusion aggregated device (AD) setup on EP ports. PR1405495
- On QFX5120, in a VXLAN-EVPN configuration, transition from collapsed to non-collapsed L2 or L3 gateway and vice versa needs a switch reload. PR1405956
- VXLAN transit traffic over a tagged underlay L3 interface and underlay IRB gets dropped due to a hardware limitation. PR1406282

- The ARP request might not be resolved successfully if the arp-suppression is enabled and **vlan-id-list** is configured on the spine node. PR1407059
- The Packet Forwarding Engine might get disabled unexpectedly because of a auto correctable non-fatal hardware error on QFX10002, QFX10008, and QFX10016. PR1408012
- DHCP discover packets might be dropped over a VXLAN tunnel if DHCP relay is enabled for other VXLAN or 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 an upgrade to Junos OS Release 17.3R1. PR1408380
- Restarting a 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 the maximum value under **vxlan-routing** on QFX Series platforms. PR1409949
- LLDP memory leak occur when IEEE DCBX packet is received in autonegotiation mode followed by another DCBX packet with none of **ieee_dcbx tlvs** present. PR1410239
- On QFX5100-48T and QFX5100-6Q, the error message **dc-pfe: BRCM_NH-,brcm_nh_resolve_get_nexthop(),346:Failed to find rt table** is seen. PR1410717
- Traffic loss might be observed after VXLAN configuration change PR1411858
- The spfe on a satellite device in a Junos fusion setup might crash and it might cause the satellite device to go offline. PR1412279
- On QFX Series platform, PEM alarm for backup FPC will be remained on master FPC though backup FPC is detached from Virtual Chassis. PR1412429
- The Junos OS device acting as the PCC might reject PCUpdate or PCCreate message if there is a metric type other than type 2. PR1412659
- On the QFX5000 line of switches, the EVPN-VXLAN multicast next-hop limit is 4000. PR1414213
- Virtual Chassis ports using DAC might not establish a link on QFX5200. PR1414492
- DC output information is missing in the **show chassis environment pem** output for whitebox. PR1414703
- VXLAN encapsulation next hop (VENH) does not get installed during BGP flapping or when routing is restarted. PR1415450
- FEC change from FEC91 to NONE does not taked effect on 100-Gigabit Ethernet interfaces with QSFP-100GBASE-SR4 optics. PR1416376
- Two instances of Junos OS are running after an upgrade to Junos OS Release 18.1R3-S3.7. PR1416585
- In Junos OS Release 18.1R3-S3, restarting routing on spine devices leads to the dcpfe generating a core file at **nh_composite_change**. PR1416925

- Rebooting QFX5200-48Y using **request system reboot** does not take physical links offline immediately. PR1419465
- During QFX5120-48Y or QFX5120-32C power cycling tests, 100-Gigabit PSM4 optics connected ports went down randomly PR1419826
- An interface might go to down state on QFX10000 and PTX10000 platforms. PR1421075
- On QFX5120-32C, DHCP binding on the client might fail when the QFX5120-32C acts as the DHCP server. This is seen only for channelized ports. PR1421110
- Fusion: ETS configuration is not applied on non-cascade ports when the AD is rebooted. PR1421429
- BFD might get stuck in slow mode on QFX10002/QFX10008/QFX100016 platform PR1422789
- QFX5100-48T 10G interface might be autonegotiated at 1-Gbps speed instead of 10Gbps. PR1422958
- The interface cannot come up when the remote-connected interface only supports 100M in QFX5100 Virtual Chassis setup. PR1423171
- ON QFX5120-32C , BUM traffic coming over irb underlay interface gets dropped on destination vtep in PIM-based VXLAN. PR1423705
- Traffic is dropped after FPC reboot with AE member links deactivated by remote device. PR1423707
- Ping over an EVPN type-5 route to QFX10000 does not work. PR1423928
- All interfaces will be down and the dcpfe might crash if SFP-T is inserted in a QFX5210. PR1424090
- IPv6 neighbor solicitation packets for link-local addresses are dropped when passing through QFX10002-60C. PR1424244
- All interfaces creation fails after NSSU. PR1425716
- Heap memory leak might be seen on QFX10000 platforms. PR1427090
- The rpd process might generate a core file because of the improper handling of graceful restart stale routes. PR1427987
- QFX5120-48Y interface with the optic QSFP-100GBASE-ER4L does not come up in "18.3R1-S2.1" PR1428113
- On QFX Series EVPN-VXLAN, the l2ald process crashes and generates a core file when the number of hardware VXLAN IFBDS exceeds the maximum limit of 16382. PR1428936
- DHCP relay might not work in an EVPN VXLAN scenario. PR1429506
- An interface on a QFX Switches does not come up after the transceiver is replaced with one having different speed. PR1430115
- In collapsed VGA4 script ping on shared ESI R6 to R7 IRB address fails. PR1430327
- On QFX Series switches, the **Validation of metadata files failed** message is seen on the hypervisor. PR1431111

- QFX5110 SFP-T: All ingress traffic is dropped on 100M fixed speed port with no-autonegotiation. PR1431885
- Transit DHCPv6 packets might be dropped on QFX5000 platforms PR1436415
- On QFX5110, QFX5200, QFX5210, there is no jnxFruOK SNMP trap message when only the power cable is disconnected and connected back. PR1437709

Interfaces and Chassis

- Constant dcpfe process crash might be seen when an unsupported GRE interface configuration is used. PR1369757
- Changing the value of **mac-table-size** to default might lead all FPCs to reboot. PR1386768
- Missing mandatory ICCP configuration statement **redundancy-group-id-list** produces a misleading error message. PR1402606
- 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 to be silently dropped and discarded. PR1417729

Layer 2 Features

- VXLAN next hop entry leak issue is seen on EX4600 and QFX5000 platforms. PR1387757
- With **IGMP snooping** enabled on the leaf switches, multicast traffic is forwarded to VLAN/VNI that does not have an active receiver. PR1388888
- On QFX Series, the error message **Failed with error (-7) while deleting the trunk 1 on the device 0** is seen. PR1393276
- On QFX5000 platforms, symmetric hashing can be done though it can not be enabled and stored in the Junos OS configuration. PR1397229
- On EVPN-VXLAN, dcpfe is restarted 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
- On QFX5000, dcpfe process crash might be observed during restart of Packet Forwarding Engine on a system with scaled EVPN-VXLAN configuration. PR1403305
- On QFX Series EVPN-VXLAN, the unicast IPv6 NS message gets flooded on L3GW. Both IPv4 and IPv6 traffic drops on L2SW. 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 impact OSPF hello packets flooding after a VTEP leaf reboot. PR1406502
- QFX5110VC generates DDOS messages of different protocols on inserting a 1G/10G SFP or forming VCP connection PR1410649
- With **arp-suppression** enabled, the QFX5000 might not forward IPv6 router solicitations or advertisement packets. PR1414496

Network Management and Monitoring

- The chassisd might crash and restart after the AGENTX session between master(snmpd) and sub-agent timeout. PR1396967
- Log files might not get compressed during the upgrade. PR1414303

Routing Protocols

- BUM packets might get looped if EVPN multihoming interface flaps PR1387063
- EVPN-VXLAN NON-COLLAPSED: AUTONEG errors and flush operation failed errors are seen after the device is power cycled. PR1394866
- On QFX5110 and QFX5200, EVPN-VXLAN NON-COLLAPSED: dcfpe generates a core file at **brcm_pkt_tx_flush, l2alm_mac_ip_timer_handle_expiry_event_loc**, after a random event. PR1397205
- On QFX5110, firewall filter applied on a VXLAN mapped VLAN is not supported in a EVPN-VXLAN scenario. PR1398237
- The rpd generates a core file and inappropriate route selection might be seen when L2VPN is used PR1398685
- The FPC/dcpfe process might crash because of interface flapping. PR1408428
- Host-generated ICMPv6 RA packets might be dropped on the backup member of VC if **IGMP-snooping** is configured. PR1413543
- The QFX Series switch might not install all IRB MAC addresses in the initialization PR1416025
- After an IRB logical interface is deleted, the MAC entry for the IRB interface is deleted for the IRB hardware address, and packets destined to other IRB logical interfaces where MAC is not configured are impacted. PR1424284

Spanning Tree Protocols

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

Resolved Issues: 18.4R1

EVPN

- The QFX10000 might drop transited traffic coming from the MPLS network to VXLAN-EVPN. PR1360159
- 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

General Routing

- After clearing the QFX5100 is treating 40G AOC uplink as 4x10g breakout with auto-channelization enabled. PR1317872
- 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 a Virtual Chassis reboot. PR1337340
- QFX5100 40G port has an interoperability issue with some other vendors. PR1349664
- 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
- The QFX5120-48Y cannot match on user-vlan-id for tunnel terminated packets. PR1358669
- On the QFX10000 line of switches, 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
- VME interface might be unreachable after link flap of em0 on master FPC. PR1362437
- Traffic might not be forwarded when the member link of the aggregated Ethernet interface is added or deleted. PR1362653
- A 1G interface might stop working when autonegotiation is off by default. PR1362977
- The following log messages are seen: **kernel: 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 with aggregated Ethernet interfaces 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
- QFX5110-5100 VCF / 1G link does not come up. PR1366218
- The tagged traffic is dropped in the untagged EVPN/VXLAN scenario. PR1366336
- 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 Cisco might go down. PR1368913
- The **commit** or **commit check** might fail due to 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
- The first 2 characters out of 14 of AS7816-64 serial number are truncated. PR1371126
- For Junos OS Release 18.1R1 and earlier releases, the USB image installation on QFX5210-64C, AMI bios upgrade needs to be done. PR1371199
- On the QFX10000 line of switches, before the Junos OS Release 17.3R3 code, the maximum number of ESI logical interfaces was 4000 in the Packet Forwarding Engine. PR1371414
- On QFX5100, the IPv6 routed packet will be transmitted though VRRP state 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 a QFX10000 line switch and a non-Juniper Network device. PR1374321
- On QFX5110, Ethernet switching flood group shows incorrect information. PR1374436
- Only the loopback interface is supported under VRF routing instances. PR1375130
- Packet Forwarding Engine wedge might be observed if there are interfaces going to down state. PR1376366
- The same address family (subnet logical interface or IRB logical interface, but not both) needs to be configured for establishing VTEPs. PR1376996
- The autonegotiation interface might go down if the opposite device supports only 10/100M autonegotiation. PR1377298

- Debug logs are printed as error logs in **/var/log/messages**. **expr_nh_flabel_check_overwrite: Caller nh_id params** message is classified as error log when it should be LOG_INFO. PR1377447
- Deleting an IRB interface might affect other IRB interfaces if the same custom MAC address is configured. PR1379002
- LOC and Diag system LED's on the front panel are not defined yet. PR1380459
- L3VPN traffic might be dropped due to one core-facing interface being down. PR1380783
- A QFX5xxx Packet Forwarding Engine might show DISCARD next-hop for overlay-bgp-lo0-ip in a spine-and-leaf topology. PR1380795
- Virtual Chassis master is copying **/var/db/ovsdatabase** to backup every 10 seconds, which causes a high write IO and shortens the SSD lifetime in Open vSwitch Database (OVSDB) environment. PR1381888
- EVPN-VXLAN ARP/NDP proxy is not working. PR1382483
- The Packet Forwarding Engine might crash if the GRE destination IP is resolved over another GRE tunnel. PR1382727
- The functionality under the license "JUNOS-FP-C2" might take effect even it does not get installed properly. PR1383274
- The 'force-host' upgrade is required for QFX5110-48S-4C in Junos OS Release 18.4 if the PTP over IPv6 G.8275.2 feature configured. PR1384073
- The Layer 3 interface might stop pinging directly connected link address after deleting Layer 2 on a physical interface. PR1384144
- On QFX5110 platforms, SFPP-10G-DT-ZRC2 and SFPP-10G-CT50-ZR transceivers might not be tunable and remain 1550.10nm by default in the hardware. PR1384524
- Port-mirroring-instance or analyzer-based mirroring does not work with input as VLAN ingress when VLAN is mapped to VXLAN. PR1384732
- 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
- CPSM daemon memory leak occurs in VMHOST. PR1387903
- On the QFX10000 line of switches, MAC learning might stop working on some LAG interfaces after frequent MAC moves. PR1389411
- FPC might crash on QFX5100 platforms in a large-scale scenario. PR1389872
- The vmcore might be seen when routing changes are made on the peer spine in an EVPN-VXLAN scenario. PR1390573
- The smid core file is seen during sanity script execution on QFX5100. PR1391909
- The l2ald core file is seen when a Layer 2 learning traceoptions were enabled. PR1394380

- DRAM and buffer utilization fields are not correct for QFX10000 platforms. PR1394978
- DOT1XD core file is found at pnac_bd_create pnac_bdm_handler knl_async_receive_and_process. PR1395384
- On QFX5110 Virtual Chassis, after Routing Engine switchover, LACP will be brought down on the peer device and never recover automatically. PR1395943
- The Juniper Extension Toolkit (JET) or Junos Telemetry Interface (JTI) is not used, because of a bug in the GRPC stack which is used by jsd and na-grpcd daemons. PR1398398

Interfaces and Chassis

- Stating in Junos OS 17.2R1, on QFX Series products, the CLI allows you to configure more logical interfaces than the limit of 2048 logical interfaces on the LAG interface. PR1361689
- On QFX5200 MC-LAG parse_remove_ifl_from_routing_inst() **ERROR : No route inst on et-0/0/16.16386**, error is seen after restarting l2cpd daemon. PR1373927

Layer 2 Features

- On QFX5100, storm control profile is missing for interfaces in hardware. PR1354889
- LACP packets are getting dropped with **native-vlan-id** configured after reboot. PR1361054
- QFX5000 the 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 interface, the LACP session to the multihomed server goes down. PR1369424
- DHCP discover packets might be dropped if VXLAN is configured. PR1377521
- Packets might be dropped on AD in a Junos Fusion Data Center environment. PR1377841
- The dcpfe process might crash while changing MTU of physical ports for GRE. PR1384517
- The LACP might be in detached state when deleting **native-vlan-id** on aggregated Ethernet interface with **flexible-vlan-tagging** configured. PR1385409
- On QFX5000 line switches, if EVPN-TYPE 5 routes are present, when doing "restart routing" or a BGP session to a neighbor device flaps, the dcpfe core file might be seen. PR1387360
- On QFX5000, EVPN-VXLAN failed 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 the Virtual Chassis master flaps. PR1389695
- Cisco Discovery Protocol (CDP) packets are not forwarded by QFX10000 line switches. PR1389829

MPLS

- LSP might not be established properly between QFX5000 line switch and other devices. PR1351055
- NO-propogate-TTL acts on MPLS swap operation. PR1366804
- LSP with auto-bandwidth enabled goes down during HMC error condition. PR1374102
- LSP "statistics" and "auto-bandwidth" functionality might not take effect with single-hop LSPs. PR1390445

Network Management and Monitoring

- For QFX5110, the returned SNMP values of module temperature-HighAlarmThreshold, LowAlarmThreshold, and HighWarningThreshold are not as same as the one shown in the CLI. PR1369030

Platform and Infrastructure

- When chassis control restart is done with aggregated Ethernet and CoS rewrite configuration, the **Platform failed to bind rewrite** messages might be seen in the syslog. PR1315437
- When Junos OS next hop index allocation fails, the private index space get exhausted through the incoming ARP requests to the management interface. PR1360039
- Forwarding is broken after adding protocol EVPN **extended-vlan-id**. PR1368802
- Traffic is silently dropped or discarded with indirect next hop and load balancing. PR1376057
- LSI binding is missing upon nd6 entry refresh after Layer 2 logical interface flap. PR1380590
- IRB interface does not turn down when master of Virtual Chassis is rebooted or stopped. PR1381272

Routing Protocols

- On QFX5100 platforms, the parity errors in Layer 3 IPv4 table in the Packet Forwarding Engine memory might cause traffic to be silently dropped and discarded. PR1364657
- On QFX5120 platforms, the command output for the configuration statement **show pfe route summary hw** shows different scale values for the IPv4 and IPv6 lpm routes rather than the supported scale. PR1366579
- The dcpfe might crash and all interfaces flap. PR1369011
- When **ecmp-resilient-hash** is configured for the existing ECMP route, the update to the next hop in hardware fails. PR1387713
- The **show evpn igmp-snooping database extensive** command output needs to be modified as per the SMET functionality. PR1391406

User Interface and Configuration

- Adding or deleting the VLAN member starting with a VLAN-ID number might cause many errors. PR1362535

SEE ALSO

New and Changed Features | **222**

Changes in Behavior and Syntax | **233**

Known Behavior | **237**

Known Issues | **240**

Documentation Updates | **262**

Migration, Upgrade, and Downgrade Instructions | **263**

Product Compatibility | **277**

Documentation Updates

There are no documentation errata or changes for the QFX Series switches in Junos OS Release 18.4R2.

SEE ALSO

New and Changed Features | **222**

Changes in Behavior and Syntax | **233**

Known Behavior | **237**

Known Issues | **240**

Resolved Issues | **249**

Migration, Upgrade, and Downgrade Instructions | **263**

Product Compatibility | **277**

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 **18.4** 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 18.4 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-x86-64-18.4-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 18.4 **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-x86-64-18.4R2.9.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-x86-64-18.4R2.9.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 18.4R1.

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-x86-64-18.4R2.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-x86-64-18.4R2.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. Disable nonstop-routing (if enabled):

   ```
   user@switch# delete routing-options nonstop-routing
   ```

6. Save the configuration change on both Routing Engines:

   ```
   user@switch# commit synchronize
   ```

7. 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.

8. 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.

9. 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-x86-64-18.4R2.n-secure-signed.tgz
   ```

   For more information about the **request system software add** command, see the CLI Explorer.

10. 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.

11. 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.

12. 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.

13. 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.

14. 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)
    ```

15. 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-x86-64-18.4R2.n-secure-signed.tgz
    ```

    For more information about the **request system software add** command, see the CLI Explorer.

16. 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.

17. Log in and issue the **show version** command to verify the version of the software installed.

18. 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.

19. 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 274
- Upgrading the Software Using Unified ISSU on page 274

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-x86-64-18.4R1.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-x86-64-18.4R2.n-secure-signed.tgz ...
     Install jinstall-host-qfx-5-f-x86-64-18.4R2.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 | 222
- Changes in Behavior and Syntax | 233
- Known Behavior | 237
- Known Issues | 240
- Resolved Issues | 249
- Documentation Updates | 262
- Product Compatibility | 277

Product Compatibility

IN THIS SECTION

- Hardware Compatibility | 277

Hardware Compatibility

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

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

Hardware Compatibility Tool

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

SEE ALSO

Junos OS Release Notes for SRX Series

IN THIS SECTION

These release notes accompany Junos OS Release 18.4R2 for the SRX Series. They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

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

> **NOTE:** The SRX5K-SPC3 Services Processing Card was introduced in Junos OS Service Release 18.2R1-S1 and is supported in all subsequent Junos OS Releases. The features and functionalities of the SRX5K-SPC3 card are supported in Junos OS Release 18.4R1. Going forward, future improvements for SRX5K-SPC3 will be included in upcoming Junos OS Maintenance Releases.

New and Changed Features

IN THIS SECTION

This section describes the new features and enhancements to existing features in Junos OS Release 18.4R2 for the SRX Series devices.

Junos OS Release 18.4R2 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-D150. Security features delivered in Junos OS for SRX Series “X” releases after 15.1X49-D150 are not available in 18.4 releases.

Release 18.4R2-S1 New and Changed Features

Chassis Clustering

- **Increase in the maximum number of child links (SRX4600)**—Starting in Junos OS Release 18.4R2-S1, you can configure eight child link interfaces in a redundant ethernet bundle on each node of the chassis cluster.

- **Dedicated fabric ports support (SRX4600)**—Starting in Junos OS Release 18.4R2-S1, you can use the built-in dedicated fabric ports as fabric link ports in chassis cluster mode.

  [See Understanding Chassis Cluster Slot Numbering and Physical Port and Logical Interface Naming, SRX Series Chassis Cluster Configuration Overview, and Chassis Cluster Control Plane Interfaces.]

Release 18.4R2 New and Changed Features

There are no new features in Junos OS Release 18.4R2 for the SRX Series devices.

Release 18.4R1 New and Changed Features

Application Security

- **CLI enhancements to support J-Web (SRX Series and vSRX)**—Starting in Junos OS Release 18.4R1, the **show service application-identification** command is enhanced to display applications and application group details in J-Web.

  The **show service application-identification** command used with the new **entries** option provides the following functionality:

  - Alphabetical list application and application group details.
  - Pagination support to limit the number of entries in output.
  - Display of details in a sorted order.
  - Using filters on output columns to search applications easily.

  [See show services application-identification entries.]

- **SSL decryption port mirroring (SRX Series and vSRX)**—Junos OS Release 18.4R1 introduces SSL decryption mirroring for SSL forward and reverse proxy. SSL decryption mirroring enables you to forward a copy of SSL decrypted traffic to a configured mirror port on a server that is acting as a traffic collection tool.

  To use the decryption mirroring feature, configure the mirror interface and the MAC address of the port in the SSL proxy profile, and apply the SSL proxy profile as the application service in the security policy. Traffic matching the policy rule is decrypted, and a copy of SSL-decrypted traffic is forwarded to the configured mirror port.

  [See SSL Proxy.]

- **Application path selection based on link preference and priority (SRX300, SRX320, SRX340, SRX345, SRX550M, SRX1500, SRX4100 SRX4200, and vSRX)**—Starting in Junos OS Release 18.4R1, you can configure Application Quality of Experience (AppQoE) to select an application path based on the link priority and the link type when multiple links are available.

  For application path selection, a list of paths to a specific destination, which meets SLA requirements, is made available. From the list, AppQoE selects a path that matches the configured link preference.

Paths are WAN links used for forwarding application traffic. You can select an MPLS or Internet link as the preferred path, and assign a priority from the range 1-255 (value of 1 indicates highest priority).

[See Application Quality of Experience.]

- **Schedulers support for APBR (SRX Series and vSRX)**—Starting in Junos OS Release 18.4R1, support for configuring policy schedulers for an advanced policy-based routing (APBR) policy is available. Using a policy scheduler, you can schedule APBR policy execution at a specified time and enforce the policy for a specified duration.

  To use a scheduler for an APBR policy, you must create a scheduler and refer to scheduler in your APBR policy configuration. The policy scheduler activates and deactivates a policy according to the scheduled time. When the scheduler times out, the associated policy is deactivated.

  [See Advanced Policy-Based Routing.]

Chassis Cluster

- **Chassis cluster resiliency (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.4R1, a three-layered model is introduced to detect software and hardware failures that impact chassis cluster performance. Flapping of em0 and control path software or hardware failures are detected and state transitions and failovers are triggered using this model. Following are the three layers:
  - **Layer 1** : Identifies and detects the components that are causing the failures.
  - **Layer 2** : Detects the failures that are not detected by Layer1.
  - **Layer 3** : Shares the health information of the system between the two nodes over control and fabric links.

  The **set chassis cluster health-monitoring** command is introduced to enable monitoring the health of chassis cluster.

  [See Chassis Cluster Resiliency.]

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 18.4R1, 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.]

General Packet Radio Service (GPRS)

- **IPv6 support on GTP (SRX1500, SRX4100, SRX4200, SRX4600, SRX4800, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 18.4R1, GPRS tunneling protocol (GTP) traffic security inspection

is supported on IPv6 addresses along with existing IPv4 support. With this enhancement, a GTP tunnel using either IPv4 and IPv6 addresses is established for individual user endpoints (UEs) between a Serving GPRS Support Node (SGSN) in 3G or a Service Gateway (S-GW) and a Gateway GPRS Support Node (GGSN) in 3G or a PDN Gateway (P-GW) in 4G.

[See GPRS Overview.]

- **Enhancements to GTP-C Tunnel (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.4R1, the GTP-C tunnel is enhanced to support tunnel-based session distribution to speed up the tunnel setup process and load-balance the sessions between the SPUs. The GTP-C tunnels and the GTP-C tunnel sessions are distributed by the SGSN tunnel endpoint identifier (TEID) of the tunnel. Use the **set security forwarding-process application-services enable-gtpu-distribution** command to enable the tunnel-based session distribution where the GTP-C traffic of different tunnels is spread across different SPUs.

  [See GPRS Overview.]

Interfaces and Chassis

- **Support for up and down delay timers on reth interfaces (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.4R1, you can configure up and down delay timers for redundant Ethernet (reth) interfaces. The delay timers keep the reth interfaces up or down, respectively, to prevent the routing protocols from reconverging and to avoid loss of traffic during a crash or when links flap.

  On SRX series devices, the default delay timer for down hold-time is 11 seconds, and the default delay timer for up hold-time is 0 seconds. To configure the timers, include the **reth 1 hold-time down *timer*** and **reth 1 hold-time up *timer*** statements at the **[edit interfaces]** hierarchy level.

  [See hold-time (Redundant Ethernet Interfaces)].

- **Half-duplex link support (SRX340 and SRX345)**—Starting in Junos OS release 18.4R1, half-duplex mode is supported on SRX340 and SRX345 devices. Half duplex enables bidirectional communication, but signals can flow in only one direction at a time. Full-duplex communication means that both ends of the communication can send and receive signals at the same time. By default, half duplex is configured. If the link partner is set to autonegotiate the link, then the link is autonegotiated to full duplex or half duplex. If the link is not set to autonegotiation, then the link defaults to half duplex unless the interface is explicitly configured for full duplex.

  [See link-mode.]

Intrusion Detection and Protection (IDP)

- **Support for custom time bindings in a time-binding custom attack (SRX Series)**—Starting in Junos OS Release 18.4R1, you can configure the maximum time interval between any two instances of a time-binding custom attack. The range for the maximum time interval is 0 minutes and 0 seconds through 60 minutes and 0 seconds. In Junos OS releases before 18.4R1, the maximum time interval between any two instances of a time-binding attack is 60 seconds.

The **interval *time-interval*** statement is introduced at the **[edit security idp custom-attack *attack-name* time-binding]** hierarchy to configure a custom time-binding.

[See Understanding Custom Attack Objects and time-binding.]

- **User visibility improvements for IDP attacks within an IDP Policy (SRX Series and vSRX)**—Starting in Junos OS Release 18.4R1, you can view and validate the complete set of attacks that are configured for an IDP policy (predefined, dynamic, and custom attacks).

  Use the **show security idp attack attack-list policy *policy-name*** command to view the attacks that are configured for an IDP policy.

  [See show security idp attack attack-list policy.]

- **IDP policy rematch (SRX Series)**—Starting in Junos OS Release 18.4R1, when a new IDP policy is loaded, the existing sessions are inspected using the newly loaded policy and are not ignored for IDP processing.

  [See IDP Policies Overview.]

Logical Systems and Tenant Systems

- Starting in Junos OS Release 18.4R1, the following features that are supported on the logical systems are now extended to tenant systems:

  - **Dynamic address support for tenant systems (SRX Series)**—Starting in Junos OS Release 18.4R1, the tenant system user can create dynamic address entries within a tenant system. A dynamic address entry contains IP ranges extracted from external sources. The security policies use the dynamic address in the **source-address** or **destination-address** field. The tenant system administrator can view the dynamic address information, including name, feeds, properties, and number of IPv4 and IPv6 entries for tenant systems, by using the **show security dynamic-address** command.

    [See Security Policies for Tenant Systems.]

  - **DHCP support for tenant systems (SRX Series)**—Starting in Junos OS Release 18.4R1, DHCP provides support for DHCP clients, DHCP relay agents, and IPv6 dynamic servers for prefix delegation for tenant systems. The DHCP relay agent operates as the interface between DHCP clients and IPv6 dynamic server for tenant systems, and also relays DHCP messages between DHCP clients and DHCP servers on different IP address networks.

    [See DHCP for Tenant Systems.]

  - **SRX5K-SPC3 card support for tenant systems (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.4R1, support for the SRX5K-SPC3 services processing card is introduced for tenant systems.

    [See Tenant Systems Overview.]

  - **Application firewall support on tenant systems (SRX Series)**—Starting in Junos OS Release 18.4R1, the tenant system administrator can configure the application firewall profile, trace options, and resources **appfw-rule-set** and **appfw-rule** in a tenant system. The application firewall rules can be

reordered using the command **insert tenants *tenant-id* security application-firewall rule-sets *ruleset-name* rule *rule-name1* after rule *rule-name2***.

Application firewall is a group of fine-grained application control policies to allow or deny the traffic based on the dynamic application name or the group names. It enhances security policy creation and enforcement based on the applications rather than traditional port and protocol analysis.

[See Application Firewall Services for Tenant Systems.]

- **Interfaces support enhancement on tenant systems (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 18.4R1, support for interfaces is enhanced on tenants systems with the following changes:
  - You can configure an interface in the tenant system similar to how you configure an interface in a logical system.
  - All types of interfaces that can be configured in a logical system can also be configured in a tenant system.
  - All the interfaces that are configured in a tenant system are associated with the routing instance configured for that tenant system.

  [See Tenant Systems Overview.]

Network Management and Monitoring

- **RPM probe enhancement (SRX Series)**—Starting in Junos OS Release 18.4R1, if the result of a probe or test exceeds the packet loss threshold, the real-time performance monitoring (RPM) test probe is marked as failed. The test probe also fails when the round-trip time (RTT) exceeds the configured threshold ranges from 0 through 60000000 ms. As a result, the device generates an SNMP notification (trap) and marks the RPM test as failed.
  RPM allows you to perform service-level monitoring. When RPM is configured on a device, the device calculates network performance based on packet response time, jitter, and packet loss.

  [See RPM Overview.]

- **SNMP support for monitoring the 4G LTE Mini-Physical Interface Module (Mini-PIM) status (SRX300, SRX320, SRX340, SRX345, and SRX550M)**—Starting in Junos OS Release 18.4R1, you can monitor 4G LTE Mini-PIM status by using SNMP remote network management.

  You can use the following commands to monitor the 4G LTE Mini-PIM status:

  show snmp mib walk ascii jnxWirelessWANNetworkInfoTable

  show snmp mib walk ascii jnxWirelessWANFirmwareInfoTable

  In previous releases, the **show modem wireless network interface *interface-name*** and **show modem wireless firmware interface *interface-name*** commands are used to check the 4G LTE Mini-PIM status.

  [See Enterprise-Specific SNMP MIBs Supported by Junos OS.]

Routing Protocols

- **ARP policer support to protect Routing Engine (SRX Series)**—Starting in Junos OS Release 18.4R1, you can apply policers on Address Resolution Protocol (ARP) traffic on SRX Series devices. You can configure rate limiting for the policer by specifying the bandwidth and the burst-size limit. Packets exceeding the policer limits are discarded.
  The traffic to the Routing Engine is controlled by applying the policer on ARP traffic. Using policers helps prevent network congestion caused by broadcast storms.

  [See ARP Policer Overview.]

Security

- **New operational commands for security policy configuration (SRX Series and vSRX)**—Starting in Junos OS Release 18.4R1, the following operational commands are introduced:
  - **show security policies information**
  - **show security policies checksum**
  - **request security policies check**
  - **request security policies resync**

  The **show security policies information** command provides detailed information about the policies configured on SRX Series devices and on vSRX. The **show security policies checksum**, **request security policies check**, and **request security policies resync** commands are used to synchronize security policies between the Routing Engine and the Packet Forwarding Engine.

  [See show security policies information, show security policies checksum, request security policies check, and request security policies resync.]

- **URL category-based security with unified policies (SRX Series)**—Starting from Junos OS Release 18.4R1, the unified policies feature is enhanced to include URL categories as match criteria for traffic flowing through the firewall. The URL category for Web filtering enables redirecting the traffic based on configured URL Category policy for further processing on the SRX Series devices. URL categories can be configured for unified policies with or without **dynamic-application** applied.

  A URL category can be configured as **url-category any** and **url-category none**. If **url-category** is not configured, the functionality is similar to **url-category none**.

  [See Configuring Unified Security Policies.]

Juniper Sky Advanced Threat Prevention

- **Juniper Sky ATP Logical Domain Support**—Starting in Junos OS 18.4, SRX Series devices support logical domains for anti-malware and security-intelligence policies. When you associate a logical domain with a realm in Juniper Sky ATP, that domain receives the threat management features configured for the realm. The SRX Series device will then perform policy enforcement based on logical domain and the associated Juniper Sky ATP realm. See *Tenant Systems: Security-Intelligence and Anti-Malware Policies* in the Juniper Sky Advanced Threat Prevention Administration Guide for details.

Software Licensing

- **Support to stop log messages on throughput overuse (SRX4100)**—Starting with Junos OS Release 18.4R1, the enhanced performance upgrade license is required to stop the log messages that are generated if the Internet mix (IMIX) throughput exceeds 20 Gbps and 7 Mpps on the SRX4100 device.

  [See Log File Sample Content.]

UTM

- **Avira scan engine support on antivirus module (SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 18.4R1, SRX Series devices support an on-device antivirus scan engine. The on-device scan engine Avira scans the data by accessing the virus pattern database. The antivirus scan engine is provided as a UTM module that you can download and install on your SRX Series device either manually (using the **request security utm anti-virus avira-engine** command) or by using the Internet to connect to a Juniper Networks-hosted URL or a user-hosted URL.

  [See On-Device Antivirus Scan Engine.]

VPN

- **Port-mirrored traffic support on an IPsec interface (SRX Series)**—Starting in Junos OS Release 18.4R1, if the output X2 interface of a mirror filter is configured for an st0 interface to filter traffic that you want to analyze, the packet is duplicated and encrypted by the IPsec tunnel bound to the st0 interface. This enhancement supports SRX Series devices in sending traffic mirrored from a port on an IPsec tunnel.

  [See Monitoring X2 Traffic.]

- **PowerMode IPsec (SRX4100 and SRX4200)**—Starting in Junos OS Release 18.4R1, 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 flow processing 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 18.4R1, 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.]

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 18.4R2 for the SRX Series.

Release 18.4R2 Changes in Behavior and Syntax

Application Security

- Starting in Junos OS Release 18.4R2, the SSL decryption mirroring feature is supported on redundant Ethernet (reth) interface on SRX Series devices operating in a chassis cluster.
- Starting in Junos OS Release 18.4R2, the format for setting up an automatic update of the application signature package is changed. Now you can use the YYYY-MM-DD.hh:mm format to configure the time for automatic download for application signatures. For example, following statement sets the start time as 10 AM on June 30, 2019:

  ```
  user@host# set services application-identification download automatic start-time 2019-06-30.10:00:00
  ```

  You can configure the automatic updates using the new format once you upgrade your previous Junos OS version to the supported Junos OS version (Junos OS Release 18.4R2).

Network Management and Monitoring

- **NSD Restart Failure Alarm (SRX Series)**—Starting in Junos OS Release 18.4R2, 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.]

VPN

- **Encryption algorithm (SRX Series)**—Starting in Junos OS Release 18.4R2, when AES-GCM 128-bit or AES-GCM 256-bit encryption algorithms are configured in the IPsec proposal, it is not mandatory to configure AES-GCM encryption algorithm in the corresponding IKE proposal.

  [See IPsec VPN Configuration Overview and encryption-algorithm (Security IKE).]

Release 18.4R1-S2 Changes in Behavior and Syntax

VPN

- **Encryption algorithm (SRX Series)**—Starting in Junos OS Release 18.4R1-S2, when AES-GCM 128-bit or AES-GCM 256-bit encryption algorithms are configured in the IPsec proposal, it is not mandatory to configure AES-GCM encryption algorithm in the corresponding IKE proposal.

  [See IPsec VPN Configuration Overview and encryption-algorithm (Security IKE).]

Release 18.4R1 Changes in Behavior and Syntax

Application Security

- **Changes to show security advance-policy-based-routing statistics command**—Starting from Junos OS Release 18.4R1, the **AppID Requested**, **Rule matches**, and **AppID cache hits** options are deprecated in the **show security advance-policy-based-routing statistics** command.

  The new options **App rule hit on cache hit**, **URL cat rule hit on cache hit**, **App rule hit midstream** and **URL cat rule hit midstream** are included to provide the details as shown in Table 4 on page 289:

  Table 4: show security advance-policy-based-routing statistics

Field Name	Field Description
App rule hit on cache hit	The number of times the rule with a matching entry in the application system cache (ASC) is found.
URL cat rule hit on cache hit	The number of times the rule with defined URL categories is matched.

Table 4: show security advance-policy-based-routing statistics *(continued)*

Field Name	Field Description
App rule hit midstream	The number of times a route is changed in the middle of a session because of the rule with defined application is matched.
URL cat rule hit midstream	The number of times a route is changed in the middle of a session because of the rule with defined URL categories is matched.

The modified **show security advance-policy-based-routing statistics** command provides the output as shown in the following sample:

```
user@host>  show security advance-policy-based-routing statistics

Advance Profile Based Routing statistics:
  Sessions Processed                    2
  App rule hit on cache hit             1
  URL cat rule hit on cache hit         0
  App rule hit midstream                1
  URL cat rule hit midstream            0
  Route changed on cache hits           1
  Route changed midstream               1
  Zone mismatch                         0
  Drop on zone mismatch                 0
  Next hop not found                    0
```

Chassis Cluster

- **Chassis cluster information detail operational command (SRX Series)**—Starting in Junos OS Release 18.4R1, use the **show chassis cluster information detail** command to view the chassis cluster information details for each node.

  [See show chassis cluster information.]

Flow-Based and Packet-Based Processing

- **New configuration options for flow configuration**—Starting from Junos OS 18.4R1, the **log dropped-illegal-packet** and **log dropped-icmp-packet** options are introduced under the [**edit security flow**] hierarchy-level.

  [See flow (Security Flow).]

- **Multiple collector support for J-Flow version 9 (SRX Series)**—Starting in Junos OS Release 18.4R1, for J-Flow version 9, up to four collectors can be configured under family inet and the PFE to export the flow record, flow record template, option data, and option data template packet to all configured collectors. Earlier to this release, only one collector could be configured under family inet and inet6.

Network Management and Monitoring

- **The NETCONF server omits warnings in RPC replies when the rfc-compliant statement is configured and the operation returns <ok/> (SRX Series)**—Starting in Junos OS Release 18.4R1, when you configure the **rfc-compliant** statement at the **[edit system services netconf]** hierarchy level to enforce certain behaviors by the NETCONF server, the server must not return an RPC reply that encloses both an **<rpc-error>** element and an **<ok/>** element. If the operation is successful, but the server reply would enclose one or more **<rpc-error>** elements of severity warning in addition to the **<ok/>** element, then 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 encloses both an **<rpc-error>** element of severity warning and an **<ok/>** element.

- **SSHD process authentication logs timestamp (SRX Series)**—Starting in Junos OS Release 18.4R1, the SSHD process authentication logs use only the time zone defined in the system time zone. In the earlier releases, the SSHD process authentication logs sometimes used the system time zone and the UTC time zone.

  [See Overview of Junos OS System Log Messages.]

UTM

- **security log message enhancement [SRX Series and vSRX]**— Starting in Junos OS Release 18.4R1, the security log information is enhanced to include source zone and destination zone for Web filtering, content filtering, antispam filtering, and antivirus features of UTM.

  [See Understanding Unified Policies [Unified Threat Management (UTM)].]

- **UTM default policy enhancement (SRX1500, vSRX)**—Starting with Junos OS Release 18.4R1, on SRX1500 Services Gateways and vSRX instances, UTM policies, profiles, MIME patterns, filename extensions, customer message, and protocol-command numbers of values are increased from 500 to 1500. The custom URL patterns and custom URL category values are increased from 1000 to 3000.

  [See UTM Overview.]

- **Antivirus profiles enhancement (SRX Series)**— Starting in Junos OS Release 18.4R1, you can create a common antivirus profile for different antivirus types. While you are creating a UTM policy for an antivirus profile, the UTM policy configuration page provides common antivirus profile selection fields for each supported protocol.

  In Junos OS Release 18.3R1 and earlier releases, separate antivirus profiles are created for every antivirus protocol. While you are creating a UTM policy for an antivirus profile, the UTM policy configuration page provides separate antivirus profile selection fields for every supported protocol.

  [See Full Antivirus Protection.]

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 18.4R2, 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 | 279

Known Behavior | 292

Known Issues | 294

Resolved Issues | 297

Documentation Updates | 309

Migration, Upgrade, and Downgrade Instructions | 309

Product Compatibility | 310

Known Behavior

This section contains the known behaviors, system maximums, and limitations in hardware and software in Junos OS Release 18.4R2 for the SRX Series.

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

Application Firewall

- SRX1500 with AppFW configured, the expected HTTP CPS is 60,000, which is a 14 percent drop (the expected value is 70,000). PR1339131

Flow-Based and Packet-Based Processing

- When user configures an interface to a zone under a tenant or root system, interfaces which are rent by other tenant are listed with question mark. PR1370255

J-Web

- CLI terminal is not working in Java version 1.8 due to security restriction in running applet. PR1341956

Platform and Infrastructure

- On SRX4600 devices, the USB disk is not made available to Junos OS. However, the USB disk is available for host OS (Linux) with full access. USB is still used in the booting process (install and recovery functions). PR1283618
- When a USB device is under initialization, removing the USB device will lead to USB crash. PR1332360

Unified Threat Management (UTM)

- Starting from Junos OS 18.3 release onwards, category in APBR module and based on destination IP address is supported, category classification will occur and APBR action will be taken place. UTM web filtering will provide an information about category to APBR module for the matched/received destination IP address. PR1365931
- To make APBR custom category to work, we need to create one local utm profile. As a workaround, to create one local utm profile use **set security utm feature-profile web-filtering juniper-local profile h1 category custom action permit** command. PR1366528

VPNs

- When multiple traffic-selectors are configured on a particular VPN object, IKED will ensure that at max 1 DPD probe is sent to the peer for the configured DPD interval. DPD probe will still be sent to the peer even if there is traffic flowing over one of the tunnels for the given VPN object. PR1366585
- On an existing tunnel, if the DPD values are changed, then they won't get applied until rekey for that tunnel happens. PR1375963

- Use the file created under **set security ike traceoptions file** command to check the logs. PR1381328
- VPN tunnels flap after adding or deleting a group in a clustered setup. PR1390831

SEE ALSO

Known Issues

IN THIS SECTION

This section lists the known issues in hardware and software in Junos OS Release 18.4R2 for SRX Series devices.

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

Application Security

- Automatic application identification download stops after going over the year and reboot **set services application-identification download automatic** stanza will be removed upon upgrading to the fixed OS. You need to add the configuration to initiate an automatic download after the Junos OS upgrade. PR1436265
- On all SRX Series devices with Security Intelligence (SecIntel) in a corner case, the black or white list file opening might fail because the file pointer might be null. which might result in the ipfd process stop. PR1436455

Flow-Based and Packet-Based Processing

- IDP installation fails on one node because the AppID process gets stuck. PR1336145
- With stress TCP traffic, some invalid sessions will time out over 48 hours. PR1383139
- On all SRX Series platforms, in a 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 be dropped due to reroute failure. PR1410233
- With PMI ON, IPsec encrypted statistics on the Routing Engine **show security ipsec statistics** command is not working anymore for fragment packets. PR1411486
- 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 SRX5400, SRX5600, and SRX5800 platforms, a corrupted **Juniper Message Passing Interface (JMPI)** message might be received by the Application Central Point (AppCP), which results in the flowd process stopping on AppCP. PR1430804
- Stream mode syslog messaging is not escaping the \ correctly. PR1416093
- On the SRX1500 platform, when an interface is changed from access mode to MVRP trunk port, traffic will be blocked and dynamic VLAN cannot be learned. As a workaround, reboot the device or srxpfe after configuration. PR1438153
- when llmd is rotating database, there is possibility that a reading access a NULL db at the same time, which generates core files. PR1439186
- Rest API is not working on SRX300 Series platform. PR1445545

Interfaces and Chassis

- LFM remote loopback is not working as expected. PR1428780

J-Web

- On SRX Series platforms, the root password configured at first J-Web access (Skip to J-Web feature) does not work if password length is shorter than eight characters. PR1371353
- On the SRX300 line of devices, an IPS installation failure message is displayed when uploading IPS signature package using the TAP mode quick setup wizard. This is an intermittent issue and occurs when IPS is installed immediately after the **system zeroized** command. PR1404296

Platform and Infrastructure

- On SRX4600 and SRX4800 devices configuration commit will fail if autoinstallation CLIs exist. PR1313095
- Upgrade from Junos OS release 15.1X49-D125 to Junos OS release 17.4X1 might cause multiple flowd process file generates on SRX cluster. PR1363314

VPNs

- On SRX Series devices, in case multiple traffic selectors are configured for a peer with 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. New negotiation of these traffic selectors is triggered through other mechanisms such as traffic or by peer. PR1287168
- IKE SAs are not displayed in CLI output after failover happens on a cluster node when tunnels are established in aggressive mode. PR1424077
- On SRX5000 Series devices with SPC3, sometimes IKE SA is not seen on the device when st0 binding on the VPN configuration object is changed from one interface to another (for example, st0.x to st0.y). PR1441411

SEE ALSO

New and Changed Features | **279**

Changes in Behavior and Syntax | **287**

Known Behavior | **292**

Resolved Issues | **297**

Documentation Updates | **309**

Resolved Issues

This section lists the issues fixed in hardware and software in Junos OS Release 18.4R2 for SRX Series devices.

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

Resolved Issues: 18.4R2

Application Firewall

- Fail to match permit rule in AppFW rule set. PR1404161

Application Layer Gateways (ALGs)

- DNS requests with the EDNS option might be dropped by the DNS ALG. PR1379433
- On all SRX Series platforms, SIP/FTP ALG does not work when SIP traffic with source NAT goes through the SRX Series devices. PR1398377
- H.323 voice packets might be dropped on SRX Series devices. PR1400630
- The TCP reset packet is dropped when any TCP proxy based feature and the **rst-invalidate-session** command are enabled simultaneously. PR1430685

Chassis Clustering

- The SNMP trap sends wrong information with **Manual failover**. PR1378903
- Traffic cannot pass through cross tenants after ISSU from Junos OS Release 18.3 to Junos OS Release 18.4. PR1382467
- Traffic with domain name address might fail for 3-5 minutes after RG0 failover on SRX Series platforms. PR1401925
- The flowd process stops when updating or deleting a GTP tunnel. PR1404317
- Mixed mode (SPC3 coexisting with SPC2 cards) high availability (HA) IP monitoring fails on secondary node with **secondary arp entry not found** error. PR1407056
- The SRX Series devices might be potentially overwritten with an incorrect buffer address when detailed logging is configured under the GTPv2 profile. PR1413718
- Starting with Junos OS Release 18.4, at most, 6 Packet Data Network Gateway (PGW) connections can be contained in a PDP context response; otherwise, the response will be dropped. PR1422877

- Memory leaks might be seen on the jsqlsyncd process on SRX chassis clusters. PR1424884
- RG0 failover sometimes causes FPC offline/present status. PR1428312

Flow-Based and Packet-Based Processing

- On SRX1500 devices, fan speed goes up and down continuously. PR1335523
- Application identification classification logic has been improved for NetBIOS and RPC. PR1357093
- Control traffic loss might be seen on SRX4600 platform. PR1357591
- When activating **security flow traceoptions**, the unfiltered traffic is captured. PR1367124
- SRX1500 continues to generate an alarm on fan **Tray 0 Fan 0 Spinning Degraded**. PR1367334
- The pkid process might stop after RG0 failover. PR1379348
- On SRX1500 devices, the activity LED (right LED) for 1-Gigabit Ethernet/10-Gigabit Ethernet port is not on although traffic is passing through that interface. PR1380928
- Password recovery menu is not shown 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
- Junos OS release 18.3R1 cannot be installed through TFTP in boot loader on SRX300 line of devices. PR1390858
- 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 by MAC local fault. PR1397012
- 40-Gigabit Ethernet 100-Gigabit Ethernet ports might take a long time (about 30 seconds) to link up on SRX4600 platform. PR1397210
- SRX Series devices might not strip VLAN added by native VLAN ID command. PR1397443
- SRX Series devices connection to JIMS keeps flapping, causing failover to secondary JIMS. PR1398140
- High jsd or na-grpcd CPU usage might be seen even when JET or JTI is not used. PR1398398
- On SRX4600 and SRX5000 devices, BGP packets might be dropped under high CPU usage. PR1398407
- VLAN push might not work on SRX1500. PR1398877
- Increase DAG feed scale number to 256 from 63. PR1399314
- The authd might stop when issuing the **show network-access requests pending** command during the authd restart. PR1401249

- SRX Series device cannot obtain IPv6 address through DHCPv6 when using a PPPoE interface with a logical unit number greater than 0. PR1402066
- Unable to access SRX Series platforms if the 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 SPC3, when PowerMode IPsec is enabled, the **show security flow statistics** and **show security flow session tunnel summary** commands will 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 and/or completely fail when utilizing services that are reliant on TCP proxy. PR1403412
- Transit UDP 500/4500 traffic might not pass across SRX5000 Series devices when using SPC3/SPC2. PR1403517
- The flowd process stops and all cards go offline. PR1406210
- The RG1 failover does not happen immediately when the SPC3 card crashes. PR1407064
- The flowd process might crash if the **enable-session-cache** command is configured under the SSL termination profile. PR1407330
- The kernel might crash on the secondary node when committing **set system management-instance**. PR1407938
- Memory leak occurs if AAMW is enabled. PR1409606
- Traffic might be lost and CPU might spike high if SSL proxy is enabled. PR1414467
- 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
- The reth interfaces are now supported when configuring SSL Decryption Mirroring (mirror-decrypt-traffic interface). PR1415352
- Traffic might be dropped if SOF is enabled in a chassis cluster in active/active mode. PR1415761
- The command **show security firewall-authentication jims statistics** will output statistics of both the primary JIMS server and secondary JIMS server. PR1415987
- when enabling PMI on SRX5400, SRX5600, and SRX5800 devices with SPC3 card or SRX4100, SRX4200, and SRX46000 devices, the flowd process stops when large size packets go through IPsec tunnel with the post fragment check triggered. PR1417219
- Traffic logging shows service-name junos-dhcp-server for UDP destination port 68. PR1417423
- Traffic might be lost on the SRX Series device if IPsec session affinity is configured with **ipsec-performance-acceleration**. PR1418135

- On all SRX Series devices, if the traffic-log feature is configured, logs might incorrectly display IPv4 addresses in an IPv6 format. PR1421255
- The **show security flow session session-identifier < sessID>** is not working if the session ID is bigger than 10M on SRX4600 platform. PR1423818
- Alarms triggered due to high temperature when operating within expected temperatures. PR1425807
- PIM neighbors might not come up on SRX Series chassis cluster. PR1425884
- The IPsec traffic going through SRX5000 line of devices with SPC2 cards installed causes high SPU CPU utilization. PR1427912
- SPC3: Uneven distribution of CPU with high PPS on device. PR1430721
- SRX550M running Junos OS Release 18.4R1 shows PEM 1 output failure message, whereas with Junos OS Release 15.1X49 or Junos OS Release 18.1R3.3 it does not show any alarms. PR1433577

Interfaces and Chassis

- Switching interface mode between family ethernet-switching and family inet/inet6 might cause traffic loss. PR1394850
- On SRX1500 platform, traffic is blocked on all interfaces after configuring the **interface-mac-limit** command on one interface. PR1409018

Intrusion Detection and Prevention (IDP)

- IDP might crash with the custom IDP signature. PR1390205
- Unable to configure dynamic-attack-group. PR1418754

Installation and Upgrade

- ISSU failed from Junos OS Release 18.3R1.9 to Junos OS Release 18.4R1.4. PR1405556

J-Web

- In the J-Web dashboard, the **Security Resources** widget did not display absolute values. PR1372826
- The **Security Log Event Details** window size was increased to display all relevant information about an event. PR1373357
- J-Web now supports defining SSL Proxy and redirect (block page) profiles when a policy contains dynamic applications. PR1376117
- **Threat Assessment Report** shows overlapping text and data. PR1397884
- Special character used in the **pre-shared-key** is removed silently after a commit operation on J-Web. PR1399363
- Configuring using the CLI editor in J-Web generates an mgd core file. PR1404946
- 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 an address set using the search function results in a commit error. PR1426321
- User unable to view GUI when logged in as read-only user. User is presented with an empty page after logging in. PR1428520
- On SRX Series devices, J-Web incorrectly displays port mode access for the link aggregation interfaces despite them being configured with multiple VLAN IDs and port mode trunk. PR1430414
- IRB interface is not available in zone option of J-Web. PR1431428

Logical and Tenant Systems

- Tenant system administrator can change VLAN assignment beyond the allocated tenant system. PR1422058

Multiprotocol Label Switching (MPLS)

- The rpd might restart unexpectedly when **no-cspf** is configured and lo0 is not included under the RSVP protocol. PR1366575

Network Address Translation (NAT)

- On SRX Series devices with SPC3 in mixed mode NAT SPC3 core files are generated 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 the reth interface. PR894194
- The chassisd might stops and restart after the AGENTX session timeout between master(snmpd) and subagent. PR1396967
- Partial traffic might get dropped on an existing LAG. PR1423989

Platform and Infrastructure

- High httpd utilization after reboot failover. PR1352133
- In chassis cluster redundancy group failover scenario, on SRX5600 and 5800 platforms, if the failover is caused by interface monitoring failure, the failover on PFE side (that is data plane) might be slow (example-impact on BFD session up to several seconds). PR1385521
- Memory leak might occur on the data plane during composite next-hop installation failure. PR1391074
- The flowd process might stop if there are too many IPsec tunnels. PR1392580
- The flowd process stops if it goes into a dead loop. PR1403276
- HA failed with the failure code **HW** after loading the image. PR1406029
- Session capacity of SRX340 device does not match with SRX345 device. PR1410801

- PEM 0 or PEM 1 or FAN, I2C Failure major alarm might be set and cleared multiple times. PR1413758
- HA packets might be dropped on SRX5000 line of devices with IOC3 or IOC2 cards. PR1414460
- Complete device outage might be seen when an SPU vmcore is generated. PR1417252
- Some applications might not be installed during upgrade from an earlier version that does not support FreeBSD 10 to FreeBSD 10 (based system). PR1417321
- On SRX Series device, flowd process stops might be seen. PR1417658
- Routing Engine CPU utilization is high and eventd is consuming a lot of resources. PR1418444
- On SRX4600 devices, commit failed while configuring 2047 VLAN IDs on the reth interface. PR1420685

Routing Policy and Firewall Filters

- The **show security flow session** command now fully supports the dynamic application. PR1387449
- Memory leak in nsd causes configuration change to not take effect after a commit. PR1414319
- The flowd process(responsible for traffic forwarding in SRX) stops on SRX devices while deleting a lot of policies from Junos Space. PR1419704
- A commit warning will now be presented to the user when a traditional policy is placed below a unified policy. PR1420471
- The dynamic-address summary's IP entry count does not include IP entries in root logical system. PR1422525
- One new alarm is created **NSD fails to restart because subcomponents fail**. PR1422738
- The ipfd generates a core file while scaling cases 6-1. PR1431861

Unified Threat Management (UTM)

- Whitelist and blacklist do 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 did not properly handle or suggest custom categories. PR1406512
- On SRX Series devices, when using Unified Policies and Web filtering (EWF) without SSL proxy, the Server Name Indication (SNI) might not be identified correctly and the RT_UTM logs were recorded incomplete information. PR1410981
- The device might not look up the blacklist first in the local Web filtering environment. PR1417330
- Unable to achieve better Avira AV TP on SRX4600 devices due to reaching mbuf high watermark. PR1419064
- 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
- The **custom-url-categories** are now pushed correctly to the Packet Forwarding Engine under all circumstances. PR1426189

User Interface and Configuration

- Tenant system administrator cannot view its configuration with **Empty Database** message when using **groups**. PR1422036

VPNs

- On SRX1500 device, when configuring IPsec VPN and BGP simultaneously, the kmd process might stop and generate a core file if BGP peers reach approximately 350. All of the VPN tunnels will be disconnected during the pause. PR1336235
- SPC3 **ike sa detail** output is not showing proper traffic statistics. PR1371638
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, 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 tunnels flap after adding or deleting a configuration group in edit private mode on a clustered setup. PR1400712
- Syslog is not generated when the IKE gateway rejects a duplicate IKE ID connection. PR1404985
- Idle IPsec VPN tunnels without traffic and with ongoing DPD probes can be affected during RG0 failover. PR1405515
- Not all the tunnels are deleted when the authentication algorithm in IPsec proposal is changed. PR1406020
- Traffic drops on peer due to bad SPI after first reauthentication. PR1412316
- On SRX5400, SRX5600, and SRX5800 devices with SPC3, when the SRX Series device is configured to initiate IKEv2 reauthentication when NAT traversal is active, occasionally reauthentication might fail. PR1414193
- The flowd/srxpfe process might stop when traffic selector is used for IPsec VPN. PR1418984
- Group VPN IKE security associations cannot be established before RG0 failover. PR1419341
- The **show security ike sa detail** command shows incorrect value in the **IPSec security associations** column. PR1423249
- On SRX Series devices with SPC3, SRX Series device does not send IKE delete notification to the peer if the traffic selector configuration is changed. PR1426714
- The kmd process stops and generates a core file after running the **show security ipsec traffic-selector** command. PR1428029

Resolved Issues: 18.4R1

Application Layer Gateways (ALGs)

- When the IPsec ALG is used, the IPsec tunnel payload is dropped after the IKE or IPsec tunnel reestablishment because of a session conflict. PR1372232
- If the SIP ALG is disabled, the SIP active sessions are affected. PR1373420
- Sun RPC data traffic for previously established ALG sessions might be dropped because it matches the gate that contains old interface information. PR1387895
- A flowd process might generate core files when cross-tenant ALG traffic is sent. PR1388658
- DNS requests with the EDNS (extension mechanisms for DNS) option option might be dropped by the DNS ALG. PR1379433

Chassis Cluster

- 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 half-duplex mode. PR1149904
- On an 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
- The **show chassis environment fpc #** command, which is used to display the FPC voltage, is enhanced to show the current and power consumption for an SPC3. PR1368507
- On SRX Series devices in chassis cluster, the 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 sends RTO message to the backup session to update the rerouted interface. PR1379305
- On SRX4600 devices in a chassis cluster, the FPCs go offline if the chassis cluster IDs are more than 10. PR1390202

Class of Service (CoS)

- When the **host-outbound-traffic** statement is configured in class of service (CoS), the device stops working when a corrupted packet arrives on the Packet Forwarding Engine. PR1359767

Command-Line Interface (CLI)

- The following CLI command outputs are not displayed correctly: **show usp memory segment shm data module** and **show jsf shm module**. PR1387711

Flow-Based and Packet-Based Processing

- On SRX320, SRX340, SRX340, and SRX550 devices, the rpd process stops when you configure the **auto-bandwidth** option under the MPLS label-switched path (LSP). PR1331164
- The security logs for unified policies are improved to reflect the reason for a denied or rejected session. PR1338310
- The IPsec replay error for Z-mode traffic is observed. PR1349724
- When the output interface configured in the X2 mirrored filter is down, the flowd process might stop. PR1357347
- On SRX4200 and SRX4600 devices, when the device is being rebooted or powered on, control traffic loss is observed. PR1357591
- IDP inline-tap mode is not supported and configuration for SPC3 must be disabled. PR1359591
- The syslog usage is deprecated, use the ERRMSG for relevant messages. PR1360274
- On the secondary control plane, a multicast session leak is observed for the PIM register. PR1360373
- The application layer protocol negotiation (ALPN) fails because the SSL proxy removes the ALPN extensions from the TLS packets. PR1360820
- On the SRX550M device, traffic might be duplicated and forwarded to the wrong interface. PR1362514
- The **show services application-identification statistics applications** command displays the **application-system-cache** error message. PR1363033
- On SRX Series devices, application identification (AppID) is supported for HTTP, SMTPS, POP3S, and IMAPS protocols. PR1365810
- When RG0 failover occurs, the flowd process generates core files. PR1366122
- The **request services user-identification authentication-table delete authentication-source** command output displays incorrect results. PR1366767
- On SRX Series devices, when AppQoE is enabled and the traffic starts flowing, the flowd process might stop. PR1367599
- On an SRX1500 device with Junos OS Release 15.1X49-D140, the srxpfe process might not work. PR1370900

- The device under test (DUT) sends incorrect rejection code when the destination device is not reachable. PR1371115
- The SPC3 core file size is larger than the SPC1 and SPC2 core files. PR1371447
- On SRX4100 and SRX4200 devices, the UDP IMIX throughput is decreased. PR1373019
- In chassis cluster mode with the IPsec tunnel configured, packet loss is observed when the clear-text packets are processed. PR1373161
- Using the SPC3 improves the performance of the unified policies. PR1374231
- A **summary** option for the **show system security-profile assignment** command is added to provide summary of security profile assignment for the entire device. PR1376990
- The SPC3 card might be installed on any slot except slot 0, slot 1, and slot 11. PR1378178
- On SRX Series devices working in a PIM sparse mode, and located between a first-hop router and a rendezvous point (RP), if a PIM control session is created through the PIM register stop message, only the next PIM register message can be forwarded, and after this first message, the subsequent PIM register messages (also matching the PIM control session above) are wrongly dropped. PR1378295
- When the datapath-debug capture is stopped, incorrect error message is displayed. PR1381703
- On an SRX5600 device in a chassis cluster, if respmod is enabled for ICAP, the connection with the ICAP server might reset automatically. PR1382376
- On SRX300, SRX320, SRX340, SRX345, SRX550M devices, during the path MTU discovery, the control engine does not receive the message **frag needed and DF set**. PR1389428
- The **set security flow log dropped-illegal-packet** and **set security flow log dropped-icmp-packet** CLI commands are unhidden. PR1394720
- On SRX Series devices, the active flow monitoring does not work for multiple collectors. PR1396482

Interfaces and Chassis

- The virtual IP address of the Virtual Router Redundancy Protocol (VRRP) might not respond to the host-inbound traffic. PR1371516

Intrusion Detection and Prevention (IDP)

- The IDP might not be deployed because the IDP configuration cannot be committed. PR1374079
- The unified policies configured with IDP might not inspect the arbitrary sessions, and are marked as **Not Interested** within the **show security idp counters flow** command. PR1385094

J-Web

- The PPPoE interface pp0 is not displayed on the **J-Web's Interfaces > Port** page. PR1316328
- The dynamic application configuration page in J-Web does not display application signatures in the result if the signatures are searched by category field. PR1344165

- The J-Web setup does not populate the DHCP attributes. PR1370700
- The chassis cluster image is not displayed on the J-Web dashboard. PR1382219

Logical Systems

- The logical system licenses fail to bind to the tenants or logical systems after the device is rebooted. PR1380144
- The logical system license limit is increased to three. One license is for root-logical-system traffic and the other two licenses are for the logical system and the tenant to transfer the traffic. PR1384659
- Tenant for logical system installation failed on node 1 after upgrading ISSU. PR1388336

Network Address Translation (NAT)

- Source NAT sessions might fail to be created when the **port-overloading** or the **port-overloading-factor** statement is configured. PR1370279

Network Management and Monitoring

- The **show snmp mib walk etherStatsTable** command displays incorrect results. PR1335808
- The eventd process generates core file, when the incoming system log message length is at or beyond the maximum supported size. PR1366120

Platform and Infrastructure

- On SRX1500 devices, when the power supply fails, the trap sent might contain incorrect information. PR1315937
- On SRX300, SRX320, SRX340, and SRX345 devices, you are unable to lock the USB port. PR1352104
- On SRX4100 and SRX4200 devices, the SRX Network Time Protocol (NTP) client might not stay synchronized to the NTP server and as a result the device clock often switches from NTP to local time. PR1357843
- On SRX5400, SRX5600, and SRX5800 devices, log messages are seen often when an IOC card has the same identifier as the SPC card. PR1357913
- When the secure copy protocol (SCP) fails to transfer the active configuration to an archive site, the archive site also fails. PR1359424
- On SRX4600 devices, the **show chassis fan show chassis environment** command does not display any output. PR1363645
- Packet capture feature does not work after the sampling configuration is deleted. PR1370779
- On SRX Series devices in a chassis cluster, the cold synchronization process might slow down when there are many Packet Forwarding Engines installed on the device. PR1376172
- Junos OS upgrade might fail when you use 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 the VNC application. PR1333206
- The **show security policies detail** command output is modified to improve readability, particularly for unified policies. PR1338307
- The timeout value of **junos-http** is not accurate. PR1371041
- When the **dynamic address** is referenced in the dynamic-address field and the destination IP address for the traffic is matched within this dynamic address, the policy fails to match the traffic PR1372921

Routing Protocols

- If family **iso** is enabled through the GRE over IPSec tunnel, the vFPC stops working. PR1364624

Services Applications

- When the ICAP configuration and the traffic passing through are modified, core files might be generated. PR1389600
- Clearing the TCP session might not clear the redirect objects. PR1390835

System Logs

- On SRX Series devices, the following false log message is observed. are observed: **/kernel: check_configured_tpids: < interfaces > : default tpid (0x8100) not configured. pic allows maximum of 0 tpids**. PR1373668

Unified Threat Management (UTM)

- The default actions under a Web filtering profile might not work properly. PR1365389
- When the server port is configured as 443, the displayed EWF server status is **UP**. PR1383695

VPNs

- IPsec tunnel might not work when there are concurrent IKEv2 Phase 1 SA rekeys. PR1360968
- On SRX5600 and SRX 5800 devices, during a migration from VPN to AutoVPN configuration, traffic loss is observed. PR1362317
- On SRX Series devices in a chassis cluster, when the VPN configuration size reaches an internal configuration processing chunk size, the VPN tunnels might not be configured successfully and the VPN tunnels might not come up after rebooting, upgrading, or restarting 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 the **show security ipsec next-hop-tunnels** command is run. 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

SEE ALSO

Documentation Updates

There are no errata or changes in Junos OS Release 18.4R2 for the SRX Series documentation.

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 generally occur in increments beyond three releases.

You can upgrade or downgrade to the EEOL release that occurs directly before or after the currently installed EEOL release, or to two EEOL releases before or after. For example, Junos OS Releases 15.1X49, 17.3, 17.4, 18.1, and 18.2 are EEOL releases. You can upgrade from one Junos OS Release to the next release or one release after the next release. For example, you can upgrade from Junos OS Release 15.1X49 to Release 17.3 or 17.4, Junos OS Release 17.4 to Release 18.1 or 18.2, and from Junos OS Release 18.1 to Release 18.2 or 18.3 and so on.

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

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

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

For information about ISSU, see the Chassis Cluster User Guide for Security Devices.

SEE ALSO

- New and Changed Features | **279**
- Changes in Behavior and Syntax | **287**
- Known Behavior | **292**
- Known Issues | **294**
- Resolved Issues | **297**
- Documentation Updates | **309**
- Product Compatibility | **310**

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 | **279**
- Changes in Behavior and Syntax | **287**
- Known Behavior | **292**
- Known Issues | **294**
- Resolved Issues | **297**
- Documentation Updates | **309**
- Migration, Upgrade, and Downgrade Instructions | **309**

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.

Licensing

Starting in 2019, Juniper Networks introduced a new software licensing model. The Juniper Flex Program is a framework, set of policies, and tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that have been developed at Juniper Networks over the past several years.

The major components of the framework are:

- A focus on customer segments (enterprise, service provider, and cloud) and use cases for Juniper Networks hardware and software products.
- The introduction of a common three-tiered model (standard, advanced, and premium) for all Juniper Networks software products.
- The introduction of subscription licenses and subscription portability for all Juniper Networks products, including Junos OS and Contrail.

For information on the list of supported products, see Juniper Flex Program.

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**—Determine the features supported on MX Series, PTX Series, QFX Series devices. The Juniper Networks Feature Explorer is a Web-based app that helps you to explore and compare Junos OS feature information to find the right software release and hardware platform for your network. https://pathfinder.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:

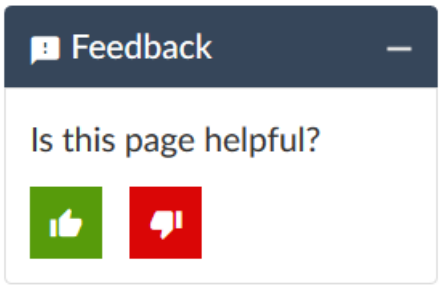

  - Click the thumbs-up icon if the information on the page was helpful to you.

- 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 Juniper Care or Partner Support Services support contract, or are covered under warranty, and need post-sales technical support, you can access our tools and resources online or open a case with JTAC.

- JTAC policies—For a complete understanding of our JTAC procedures and policies, review the *JTAC User Guide* located at https://www.juniper.net/us/en/local/pdf/resource-guides/7100059-en.pdf.
- Product warranties—For product warranty information, visit https://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://www.juniper.net/customers/support/
- Search for known bugs: https://prsearch.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://www.juniper.net/customers/csc/software/
- Search technical bulletins for relevant hardware and software notifications: https://kb.juniper.net/InfoCenter/
- Join and participate in the Juniper Networks Community Forum: https://www.juniper.net/company/communities/
- Create a service request online: https://myjuniper.juniper.net

To verify service entitlement by product serial number, use our Serial Number Entitlement (SNE) Tool: https://entitlementsearch.juniper.net/entitlementsearch/

Creating a Service Request with JTAC

You can create a service request with JTAC on the Web or by telephone.

- Visit https://myjuniper.juniper.net.
- 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, see https://support.juniper.net/support/requesting-support/.

Revision History

15 July 2021—Revision 15, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

13 January 2021—Revision 14, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

23 July 2020—Revision 13, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 May 2020—Revision 12, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 January 2020—Revision 11, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 November 2019—Revision 10, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

31 October 2019—Revision 9, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 October 2019—Revision 8, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

3 October 2019—Revision 7, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

12 September 2019—Revision 6, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

30 August 2019—Revision 5, Junos OS Release 18.4R2– 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 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

25 July 2019—Revision 3, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 July 2019—Revision 2, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

10 July 2019—Revision 1, Junos OS Release 18.4R2– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

16 May 2019—Revision 12, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

22 April 2019—Revision 11, Junos OS Release 18.4R1– 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 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

15 March 2019—Revision 9, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

14 February 2019—Revision 8, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

7 February 2019—Revision 7, Junos OS Release 18.4R1– 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 18.4R1– 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 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

17 January 2019—Revision 4, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

11 January 2019—Revision 3, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

28 December 2018—Revision 2, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.

21 December 2018—Revision 1, Junos OS Release 18.4R1– ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, and Junos Fusion.