Release Notes

Published
2023-06-01

Junos® OS Release 21.1R3 for the ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX

Juniper Networks | Engineering Simplicity

Table of Contents

Introduction

Junos OS runs on the following Juniper Networks® hardware: ACX Series, cSRX, EX Series, JRRSeries, JuniperSecureConnect,Junos Fusion Enterprise,Junos Fusion ProviderEdge, MX Series, NFXSeries, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

These release notes accompany Junos OS Release 21.1R3 for the ACX Series, cSRX Container Firewall (cSRX), EX Series, JRR Series, Juniper Secure Connect, Junos Fusion Enterprise,Junos Fusion Provider Edge, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, virtual MX Series router (vMX), Virtual Route Reflector (vRR), and vSRX Virtual Firewall (vSRX).They describe new and changed features, limitations, and known and resolved problems in the hardware and software.

Junos OS Release Notes for ACX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for ACX Series routers.

What's New in 21.1R3

There are no new features or enhancements to existing features in Junos OS Release 21.1R3 for ACX Series routers.

What's New in 21.1R2

There are no new features or enhancements to existing features for ACX Series in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for the ACX Series.

EVPN

- **Support for EVPN E-Tree service (ACX5448)**—Starting with Junos OS Release 21.1R1, you can configure an EVPN Ethernet Tree (E-Tree) service on ACX5448 routers.

  [See EVPN-ETREE Overview.]

- **Support for inter-DC connectivity over a Layer 3 network (ACX5448)**—Starting with Junos OS Release 21.1R1, you can configure the ACX5448 router to support IRB interfaces in an EVPN-MPLS network. This feature supports EVPN Type 2 (MAC/IP advertisement) and EVPN Type 5 (IP prefix) routes.

  [See EVPN with IRB Solution Overview.]

- **Support for single-active multihoming redundancy in EVPN-VPWS with flexible cross-connect support (ACX5448)**—Starting with Junos OS Release 21.1R1, you can configure the interfaces on the ACX5448 router in an Ethernet VPN–virtual private wire service (EVPN-VPWS) network with flexible cross-connect (FXC) or legacy cross-connect (non-FXC) service to support single-active multihoming redundancy for traffic that flows from customer edge devices to the core. EVPN-VPWS also supports load balancing with equal-cost multipath (ECMP) fast reroutes (FRR) on IGP and over BGP multipaths that face the core.

  [See Overview of Flexible Cross-Connect Support on VPWS with EVPN and Configuring EVPN Active-Standby Multihoming.]

- **Tunnel endpoint in the PMSI tunnel attribute field for EVPN Type 3 routes (ACX5448, EX4600, EX4650, EX9200, and QFX10002)**—Starting in Junos OS Release 21.1R1, you can set the tunnel endpoint in the provider multicast service interface (PMSI) tunnel attribute field to use the ingress router's secondary loopback address. When you configure multiple loopback IP addresses on the local provider edge (PE) router and the primary router ID is not part of the MPLS network, the remote PE router cannot set up a PMSI tunnel route back to the ingress router.

  To configure the router to use a secondary IP address that is part of the MPLS network, include the `pmsi-tunnel-endpoint`*`pmsi-tunnel-endpoint`* statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn]` hierarchy level for both EVPN and virtual-switch instance types.

  [See EVPN.]

- **Aliasing for all-active multihoming with EVPN-MPLS (ACX5448)**—Starting in Junos OS Release 21.1R1, ACX5448 routers support aliasing for EVPN-MPLS all-active multihoming with ELAN services. Aliasing enables remote provider edge (PE) devices to load balance Layer 2 traffic toward a multihomed customer edge (CE) device among the PEs that have the same EVPN segment ID (ESI) for that CE device.

  You enable aliasing when you configure the `load-balance per-packet` routing policy statement at the `[edit policy-options policy-statement]` hierarchy and `export` the policy statement at the `[edit routing-`

`options forwarding-table]` hierarchy. This feature is supported in routing instances of type `evpn` with VLAN-based and VLAN bundle services.

[See EVPN Multihoming Overview.]

MPLS

- **BGP Classful Transport planes (BGP-CT) to facilitate service mapping over colored tunnels (ACX Series, PTX Series, MX Series)**—Starting in Junos OS Release 21.1R1, you can classify colored transport tunnels (RSVP, IS-IS flexible algorithm) in your network into transport classes and map service routes over an intended transport class. You can also extend the transport tunnels to span across multiple domains (ASs or IGP areas) by using the new BGP transport address family called BGP Classful Transport (BGP CT).

  This feature lays the foundation for network slicing and allows the different domains to interoperate irrespective of the transport signaling protocols used in each domain.

  [See BGP Classful Transport Planes Overview.]

Network Management and Monitoring

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

  [See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

Routing Protocols

- **Support for configuring multiple independent IGP instances of IS-IS (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 21.1R1, you can configure and run multiple independent IGP instances of IS-IS simultaneously on a router.

> **NOTE**: Junos OS does not support configuring the same logical interface in multiple IGP instances of IS-IS.

[See How to Configure Multiple Independent IGP Instances of IS-IS.]

Segment Routing

- **Support for flexible algorithms in IS-IS for segment routing–traffic engineering (SR-TE) (ACX Series)**—Starting in Junos OS Release 21.1R1, you can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize the IGP metric and another flexible algorithm to compute a path based on the traffic engineering metric to divide the network into separate planes. This feature enables networks without a controller to configure traffic engineering and utilize the segment routing capability of a device.

  To define a flexible algorithm, include the `flex-algorithm` statement at the `[edit routing-options]` hierarchy level. To configure a device to participate in a flexible algorithm, include the `flex-algorithm` statement at the `[edit protocols isis segment routing]` hierarchy level.

  [See Understanding IS-IS Flexible Algorithm for Segment Routing.]

- **Support for flexible algorithm in OSPFv2 for segment routing traffic engineering (ACX5448, ACX710, MX204, MX104, MX480, MX960, MX10003, MX2020, and PTX10001)**—Starting in Junos OS Release 21.1R1, you can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path based on traffic engineering metric to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

  To define a flexible algorithm, include the `flex-algorithm` statement at the `[edit routing-options]` hierarchy level.

  To configure a device to participate in a flexible algorithm, include the `flex-algorithm` statement at the `[edit protocols ospf source-packet-routing]` hierarchy level.

  [See How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering.]

- **Support for strict SPF and IGP shortcut (ACX710, MX960, MX10008, MX2020, PTX5000, and PTX1000)**—Starting in Junos OS Release 21.1R1, you can configure segment routing algorithm 1 (strict SPF) and advertise its SIDs in IS-IS link-state PDU (LSPDU) and use these SIDs to create SR-TE tunnels to forward the traffic by using the shortest IGP path to reach the tunnel endpoint while avoiding loops. You can also specify a set of prefixes in the import policy, based on which the tunnel

can redirect the traffic to a certain destination. You can use algorithm 1 (strict SPF) along with algorithm 0 (default SPF) by default when Source Packet Routing in Networking (SPRING) is enabled.

[See How to Enable Strict SPF SIDs and IGP Shortcut, prefix-segment, and source-packet-routing.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for ACX Series routers.

What's Changed in Release 21.1R3

IN THIS SECTION

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Prior to this PR, the source and destination address pair had to be unique among all the FTI tunnels regardless of the tunnel encapsulation type.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

  [See request system scripts refresh-from and cert-file (Scripts).]

What's Changed in Release 21.1R2

IN THIS SECTION

General Routing

- **Changes in contextEngineID for SNMPv3 INFORMS (PTX Series, QFX Series, ACX Series, EX Series, MX Series, and SRX Series)**— Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  [See SNMP MIBs and Traps Supported by Junos OS.]

Interfaces and Chassis

- **Blocking duplicate IP detection in the same routing instance (All Junos platforms)**—Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. Refer to the table mentioned in the topic inet (interfaces). When you try to configure same IP on two logical

interfaces inside same routing instance, the commit will be blocked with the error displayed as shown below:

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24

[edit]
user@host# commit
commit complete

[edit]
user@host# set interfaces ge-0/0/2 unit 0 family inet address 2.2.2.2/24

[edit]
user@host# commit
[edit interfaces ge-0/0/2 unit 0 family inet]
  'address 2.2.2.2/24'
     identical local address found on rt_inst [default], intfs [ge-0/0/2.0 and ge-0/0/1.0],
family [inet].
error: configuration check-out failed
```

[See inet(interfaces).]

What's Changed in Release 21.1R1

IN THIS SECTION

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you

omit the parameter, the function behaves as in earlier Junos OS releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

[See invoke() Function (SLAX and XSLT).]

- **The jcs:invoke() function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The jcs:invoke() extension function supports the no-login-logout parameter in SLAX event scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding language python statement at the [edit system scripts] hierarchy level. To execute Python scripts, configure the language python3 statement at the [edit system scripts] hierarchy level to execute the scripts using Python 3.

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

Network Management and Monitoring

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to advertise third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the [edit system services netconf hello-message yang-module-capabilities] hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the [edit system services netconf netconf-monitoring netconf-state-schemas] hierarchy level.

  [See hello-message and netconf-monitoring.]

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the client-alive-interval and client-alive-count-max statements at the [edit system services netconf ssh] hierarchy level. The client-alive-interval statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The client-alive-count-max statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

[See ssh (NETCONF).]

User Interface and Configuration

Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—The Junos OS CLI exposes the `verbose` statement at the `[edit system export-format json]` hierarchy level. We changed the default format to export configuration data in JSON from `verbose` to `ietf` starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `[edit system export-format json]` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See export-format.]

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for ACX Series routers.

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

General Routing

- The PTP lock state when a source selection is changing between virtual port and PTP master indicates holdover-in-spec. PR1510880
- Currently, this metric is not met due to limitations from the microsemi servo code. PR1522796
- The issue is regarding tuning T1/T4 close to 2 way cTE . Currently, we see around -180 to 200ns cTE . As of now with the current tuning parameters , the critical parameters cTE and 2 way TE is taken care of. As the individual T1/T4 time error will not cause impact to 2 way , cTE , not fixing this. PR1527347

- The packet time error on ACX5448 chassis with g.8275.2.enh profile is exceeding class-A time error limits of max TE of 100ns. The 1pps time error is exceeding the cTE of 50ns. PR1535434
- Changing PTP profile type on ACX710 from g.8275.1 to g.8275.2 and vice versa requires a Packet Forwarding Engine reboot and clksyncd restart. PR1546614
- Configuring rib-group to import or export routes across different routing instance is not supported on DNX(ACX5448) platform. These rib-group imported/exported routes will be of type "rtbl"(rtable). To configure this rtable type routes and do route forwarding in hardware, double route lookup support is required. DNX hardware chipset does not support the double route lookup across two different routing tables. Hence rib-group route import or export across routing instance is not supported on DNX platform. PR1547078
- In PHASE ALIGNED state , clksyncd restart causes servo to move to FREERUN, but "show ptp clock " will still show GMC Id of connected master since packet exchange is established. PR1548192
- Ping might fail if the MAC address of the device is modified to a static MAC address as BCM supports only one base MAC address. PR1553472
- Unified ISSU is not supported on releases prior to Junos OS Release 20.4 to releases 20.4 and above. There is a major SDK upgrade from 6.3.2 to 6.5.16, due to which the Warm boot feature needed for unified ISSU is not supported by our vendor. PR1554915
- When ACX5448/ACX710 is configured as Two-Way Active Measurement Protocol (TWAMP) server and a client sends a start session message, the response from ACX to send "start session ack" is delayed by 10 seconds, which might cause the session to fail. PR1556829

Open Issues

IN THIS SECTION

Learn about open issues in this release for ACX Series routers.

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

General Routing

- The "ping" command on an ACX device might show variable latency values. This is expected for host-generated ICMP traffic due to the design of the PFE queue polling the packets from ASIC. PR1380145
- On ACX6360/PTX10001 router, Tx power cannot be configured by using + sign. PR1383980
- The ccc logs are not compressed after rotation. PR1398511
- Multiple RMEPs are unsupported due to which RDI issue was seen. The False Alarm RDI issue is now being tracked by PR1421845. PR1478346
- If you configure DHCP option 012 host-name in DHCP server configuration and the actual base configuration file also has the host-name in it, you might be overwriting the base configuration file's host-name with the DHCP option 012 host-name. PR1503958
- For ACX710, if the console cable is plugged in and the terminal connection is active and sending characters to the interface, the system boot might be interrupted and the ACX710 boot will be stalled at the uboot# prompt. PR1513553
- If you restart from alternate media, alarm might not be seen on ACX710, when system is restarted with recovery snapshot. This feature requires new implementation and changes in ACX710 firmware. PR1517221
- Due to BRCM KBP issue route, lookup might fail. Need to upgrade KBP to address this issue. PR1533513
- A new alarm "network-service mode mismatch between configuration and kernel setting" was introduced by PR 1514840 commit. When unified ISSU or normal code upgrade is performed from images without PR1514840 commit to images with PR1514840 commit, then the transient false alarm will be seen. PR1546002
- On ACX5448 BFD session status is in Init state, after system reboot. It is seen when we have both CFM and BFD configured on the system, due to endpoint overlap between CFM and BFD. PR1552235
- On ACX5448 platform running with Junos OS Release 19.x/20.x, it might be unable to downgrade to Junos OS release 18.x due to an issue with validation. If it is downgraded with the 'no-validate' command option, the system will be unstable after it starts, with the chassisd process restarting continuously. Please refer to TSB18096 for more details. PR1556377
- On certain ACX platforms, MAC address entries might not be deleted from the MAC table at the end of 'mac-table-aging-time' timer when there is active traffic destined to that MAC address. When the issue happens, it might reduce the number of new MAC addresses that can be learned. If the

ethernet-switching table overflows, no new MAC addresses will be learned, which might cause traffic flooding. PR1565642

- The Precision Time Protocol (PTP) clock might fail to be locking and might get stuck in an acquiring state at clock servo. PR1570310
- On MX Platforms, the MPC7E, MPC10E, MX-SPC3 and LC2103 line cards might become offline, resulting in complete loss of traffic when the device is running on FIPS mode. The CLI command 'show chassis fpc pic-status' can be used to check the status of the line cards. PR1576577
- In some scenarios with CFM, Layer 2 VPN, Layer 2 Circuit and Layer 3 VPN scaling, on reboot while receiving a CFM inline event, the session will be programmed in hardware only if the logical interface (IFL) is valid. Maximum of 4 retries can happen to program the session in hardware. But here on reboot, IFL is not valid even after 4 retries. So CFM sessions on those IFLs stuck in start state. Number of retry count has been increased to program the session in hardware. PR1602489
- This changes increase the number of family inet arp policers to 64 entries. TCAM resource shortage errors can be seen if there are more than 32 IFLs with configured arp policer. PR1630280
- When multihop BFD is configured on ACX5448, delegated BFD sessions are not coming up. PR1633395

Interfaces and Chassis

- The remote-mep-state is not as expected. PR1623960

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for ACX Series routers.

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

Resolved Issues: 21.1R3

IN THIS SECTION

General Routing

- On ACX5448 routers, the two-way time error and CTE for 1 PPS do not meet the class A metrics. PR1535434
- You might see packet buffer allocation failed messages during scaled CFM sessions with minimum DM/SLM cycle-time along with enhanced-sla-iterator. PR1574754
- PTP might get stuck and not function properly on ACX710 in a certain condition. PR1587990
- Traffic might get forwarded through the member links in down state after new member links are added to AE interface on ACX710/ACX5400. PR1589168
- ACX5448: Seeing high DMR out of sequence with iterator configuration. PR1596050
- ACX710: l2ald.core seen @ l2ald_event_process_list_id, l2ald_event_proc_all_lists, l2ald_event_periodic () at ../../../../../../src/junos/usr.sbin/l2ald/l2ald_event.c:757. PR1596908
- Traffic drop in EVPN VPWS flexible cross connect on ACX5448/710. PR1598074
- Traffic loss might be observed if "drop-profiles" is modified on ACX710/ACX5448. PR1598595
- MACsec traffic over L2circuit might not work on ACX5448 and ACX710 platforms. PR1603534
- The FPC might restart when executing the command "show firewall" on the ACX5448 platform. PR1605288
- The optics_mts_010.robot script is failing while verifying SNMP and matches CLI values. PR1605348
- ACX5448/710 platforms running DHCP relay will not process packets arriving over MPLS. PR1605854
- The Forwarding Engine Board (FEB) might crash on ACX1000, ACX1100, ACX2000, ACX2100, and ACX4000 platforms. PR1606424

- DHCP packets might not be relayed on the ACX710 and ACX5448 platforms. PR1608125
- ACX5096: output pps traffic is seen on deactivated interfaces on ACX5096. PR1608827
- The Routing protocol engine CPU is getting stuck at 100%. PR1612387
- ACX5048 places host-outbound traffic in an incorrect queue. PR1619174
- Traffic might get equally load-balanced irrespective of the scheduler configuration. PR1620137
- In Layer 3 VPN scenario with ACX5448 after multiple core link flaps the following errors might be seen dnx_nh_unilist_install_multipath: Failed to create shadow obj 0x20017ff0 for NH 766(FEC 0x2000109f) unilist nh 2097161. Error -14(No resources for operation). PR1621425

Resolved Issues: 21.1R2

IN THIS SECTION

General Routing

- The IPv6 BFD sessions flap when configured below 100 ms flaps. PR1456237
- The aggregated Ethernet interface might not come up with LFM configured after reboot. PR1526283
- In the Layer 3 VPN scenario, the CE device traffic drops on ingress PE device while resolving using default route in VRF. PR1551063
- Verifying multiple PD synchronizations with relay deletes and adds configurations. PR1554647
- The ACX5448/ACX710 router as TWAMP server delays the start session acknowledgment by 10 seconds. PR1556829
- On the ACX5448 routers, single rate three color policer does not work. PR1559665
- When an RDI is received with CCM packet, sessions are not deleted. PR1560182
- The DNX router fails to program Mcast route in BCM, when the route has pime interface as outgoing interface. PR1560914

- Inline BFD stays down with ISIS/Static clients. PR1561590
- Analyzer (Port Mirroring) might not work on ports above 20. PR1563774
- Loopback0 firewall might not take effect along with error logs. PR1566417
- Pushing more than 2 MPLS labels on ACX5448/ACX710 might not work. PR1566828
- On the ACX500 routers, service MIC does not work. PR1569103
- [interface] [interface] : acx5048 :: [IFD: traffic-input-pps not incrementing for vlantagged_flexible traffic]. PR1569763
- ACX resets tunable optics to default wavelength after upgrade/reboot. PR1570192
- On the ACX5448 routers, the untagged traffic is being incorrectly queued and marked. PR1570899
- ACX710: PEM Feed snmp trap support. PR1571368
- ACX5448: RFC2544 reflector feature could not work on a higher port. PR1571975
- The l2circuit and CFM sessions might go down with asynchronous-notification configured. PR1572722
- ARP traffic exceeding the policer limit is not discarded on ACX platforms. PR1573956
- On ACX5448/710 platforms 802.1P rewrite might not work. PR1574601
- Packets might get tagged with the default VLAN-ID and dropped at the peer under Layer 2 circuits local switching scenario. PR1574623
- ACX fails to process RSVP Path Message. PR1576585
- RLFA does not takes effect due to service label incorrectly popped. PR1577460
- "LIBCOS_COS_TVP_FC_INFO_NOT_FOUND: Forwarding-class information not specified" is seen when commiting scheduler-map under class-of-service. PR1579009
- On the ACX710 routers, continuous reboot due to configuration under auxiliary port s observed. PR1580016
- On the ACX5448 platform asynchronous-notification for 1G interface fails to work. PR1580700
- There might be a traffic drop between customer edge and provider edge devices in case of ARP resolution failure. PR1580782
- The process rpd may stuck in 100% due to race condition. PR1582226
- acx710 log jnpr-clock-recovery.log file size too small and archives rotate too quick. PR1582350

- [interface] [AE] ACX-710 :: ACX: ACX710: unexpected results observed while verifying channelised interface check with snmp mib get ifHighSpeed output. PR1583995
- ACX5448 - ACX_ASIC_PROGRAMMING_ERROR - Detection time shows the default value (6.000) instead of the configured value for single hop BFD. PR1585382
- DHCPv4 might not work on ACX710/ACX5448. PR1589135
- Traffic might get forwarded through the member links in down state after new member links are added to AE interface on ACX710/ACX5400. PR1589168
- ACX5448/710 platforms running DHCP relay will not process packets arriving over MPLS with an explicit null label. PR1590225
- Traffic is not passing through the l2circuit interface when vlan-id-range configured. PR1590969
- [eoam] [eoamtag] acx5448 :: Seeing high DMR out of sequence with iterator configuration. PR1596050
- [l2snooping] [l2snoopingtag] ACX-710 :: ACX710: l2ald.core seen @ l2ald_event_process_list_id, l2ald_event_proc_all_lists, l2ald_event_periodic () at ../../../../../src/junos/usr.sbin/l2ald/l2ald_event.c:757. PR1596908
- Traffic drop in EVPN VPWS flexible cross connect on ACX5448/710. PR1598074
- PDT-SP-METRO-After appending CoS configuration when we are changing the drop-profile fill-levels we are seeing ae interface is going down. PR1598595

Platform and Infrastructure

- Junos OS: Upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart (vmcore). PR1557881

Routing Protocols

- BGP session carrying VPNv4 prefix with IPv6 next-hop might be dropped. PR1580578

Resolved Issues: 21.1R1

IN THIS SECTION

Class of Service (CoS)

- The explicit classifier or rewrite-rule might not work as expected for a logical interface if the wildcard configuration is also applied. 1556103
- FPC might crash after committing the `show class-of-service` command. PR1568661

  </p>

Forwarding and Sampling

- VLAN-ID-based firewall match conditions might not work for the VPLS service. PR1542092

General Routing

- Memory utilization enhancement is needed. PR1481151
- The ACX1100, ACX2100, ACX2200, ACX2000, and ACX4000 routers might stop forwarding transit and control traffic. PR1508534
- On the ACX5448 routers, transit DHCP packets drop are observed. PR1517420
- On the ACX500-I router, the `show services session count` command does not work as expected. PR1520305
- PTP to 1PPS noise transfer test fails for frequency 1.985 Hz. PR1522666
- Interface does not come up with the auto-negotiation setting between the ACX1100 router and the other ACX Series routers, MX Series routers, and QFX Series switches as the other end. PR1523418
- On the ACX710 routers, PIR or CIR Hqos behaviour is inconsistent. PR1525789

- With the ACX5448 router with 1000 CFM, the CCM state does not go in the **Ok** state after loading the configuration or restarting the Packet Forwarding Engine. PR1526626
- The l2cpd process might leak memory with the aggregated Ethernet interface flap. PR1527853
- The FEC field is not displayed when the interface is down. PR1530755
- Unable to switch profile between G.8275.1 and G.8275.2. PR1533263
- Upon classifying the Layer 3 packets, DSCP is not preserved and is lost at the egress due to the limitations of a chipset. PR1535876
- The clksyncd process generates core file on Junos OS Release 20.3R1.3 image. PR1537107
- The rpd process generates a core file at `l2ckt_vc_adv_recv, l2ckt_adv_rt_flash (taskptr=0x4363b80, rtt=0x4418100, rtl=< optimized out>, data=< optimized out>, opcode=< optimized out>) at ../../../../../../../../../src/junos/usr.sbin/rpd/l2vpn/l2ckt.c:7982.` PR1537546
- Management Ethernet link down alarm is seen while verifying the system alarms in a Virtual Chassis setup. PR1538674
- On the ACX5448 router, the BGPV6LU traffic drop is observed when the node is deployed in ingress. PR1538819
- On the ACX5448 router, unexpected behavior of the `show chassis network-services` command is observed. PR1538869
- The following error message is observed while deleting the remote stream 0 0 0 0 0 0 along with feb core file at `0x00ae6484 in bcmdnx_queue_assert (queue=0xc599b60)` at `../../../../../src/pfe/common/drivers/bcmdnx/bcmdnx_sdk_ukern_layer.c`:

  ```
  Err] clksync_mimic_delete_clock_entry Unexpected error.
  ```

  PR1539953
- The announcement or synchronization interval rate range is not as expected. PR1542516
- Synchronization Ethernet goes in the **Holdover** state and comes back to the **Locked** state when the PTP configuration is deleted. PR1546681
- The ACX5448 router as transit for the BGP labeled unicast drops traffic. PR1547713
- IP addresses other than IPv4 and IPv6 must not be forwarded. PR1550748
- Multicast traffic is stopped when HQoS with multicast configurations is applied. PR1551248
- Verifying multiple PD synchronizations with relay deletes and adds configurations. PR1554647

- The ARP packets from the CE device are added with VLAN tag if the VLAN-ID is configured in the EVPN routing instance. PR1555679
- On the ACX710 router, the T-BC-P switch-over performance fails beyond the standard mask and servo moves to multiple **Holdover-in** state, **Acquiring** state, and **Holdover-out** state. PR1556087
- On the ACX5448 router, you cannot downgrade to Junos OS Release 18.4 code-base. PR1556377
- On the ACX5448 router, the unicast packets from the CE devices might be forwarded by the PE devices with an additional VLAN tag if IRB is used. PR1559084
- On the ACX5048 router, the fxpc process generates a core file on the analyzer configuration. PR1559690
- On the ACX2100 routers, laser-output-power is observed after the interface is disabled and then rebooted. PR1560501
- On the ACX5448 router, the following syslog message is reported every 30 seconds:

  ```
  ACX_DFW_CFG_FAILED: ACX Error (dfw):dnx_dfw_dyn_entry_counter_get : Entry is invalid. <url
  ```

  PR1562323
- Expected entries are not observed while verifying the physical logical interface status with the Layer 3 traffic. PR1572211
- The aggregated Ethernet interface might not come up with LFM configured after reboot. PR1526283
- Packets might drop with all the commit events with the 1G speed configured interface. PR1524614
- ACX Series routers fails to process the RSVP path message. PR1576585
- BUM traffic might drop in the VPLS instance under certain conditions. PR1531733
- Snmp mib walk for jnxSubscriber OIDs returns a general error. PR1535754
- On the ACX5448 and ACX710 routers, multicast traffic loss might occur. PR1538053
- On the ACX5448 routers, the SFP-T interface might not come up if a straight cable is used. PR1547394
- When the LACP daemon restarts, the LACP local partner system ID remains 0 in the MC-AE output. PR1560820
- Analyzer (Port Mirroring) might not work on ports above 20. PR1563774
- The DF (Designated Forwarder) might not forward traffic. PR1567752

- The ACX routers resets tunable optics to default wavelength after upgrade or reboot. PR1570192

Infrastructure

- The vme/me0 management interface does not process any incoming packets. PR1552952

Interfaces and Chassis

- The traffic loss might occur on an interface when you configure the non-related to the interface. PR1541835

Layer 2 Features

- On the ACX5448 routers, VPLS traffic statistics are not displayed when the `show vpls statistics` command is executed. PR1506981

Network Management and Monitoring

- The SNMP might not be work when ISIS is disabled under VRF. PR1527251

Routing Protocols

- The rpd memory might leak might in the BGP scenario. PR1547273

Documentation Updates

There are no errata and changes in Junos OS Release 21.1R3 for the ACX Series documentation.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for ACX Series routers. 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 https://www.juniper.net/documentation/en_US/junos/information-products/pathway-pages/software-installation-and-upgrade/software-installation-and-upgrade.html Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 1: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Junos OS Release Notes for cSRX

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for the cSRX Container Firewall, a containerized version of the SRX Series Services Gateway. 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for cSRX.

What's New in 21.1R3

There are no new features for cSRX in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features for cSRX in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for cSRX.

Authentication and Access Control

- **Configure client information to connect to the JIMS server (cSRX, SRX300, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting with Junos OS Release 21.1R1, you can configure which specifc interface, source IP address or routing instance SRX should use for connecting to a JIMS server.
  [See Configuring the Connection to an SRX Series Device.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for cSRX.

What's Changed in Release 21.1R3

There are no changes in behavior or syntax for cSRX in Junos OS Release 21.1R3.

What's Changed in Release 21.1R2

There are no changes in behavior or syntax for cSRX in Junos OS Release 21.1R2.

What's Changed in Release 21.1R1

There are no changes in behavior or syntax for cSRX in Junos OS Release 21.1R1.

Known Limitations

There are no known limitations for cSRX in Junos OS Release 21.1R3.

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

Open Issues

There are no open issues for cSRX in Junos OS Release 21.1R3.

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

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for cSRX.

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

Resolved Issues: 21.1R3

There are no resolved issues for cSRX in Junos OS Release 21.1R3.

Resolved Issues: 21.1R2

There are no resolved issues for cSRX in Junos OS Release 21.1R2.

Resolved Issues: 21.1R1

There are no resolved issues for cSRX in Junos OS Release 21.1R1.

Junos OS Release Notes for EX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for EX Series switches.

What's New in 21.1R3

There are no new features or enchancements to existing features in Junos OS Release 21.1R3 for EX Series Switches.

What's New in 21.1R2

There are no new features or enchancements to existing features in Junos OS Release 21.1R2 for EX Series Switches.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for EX Series Switches.

Hardware

- **New EX4400 switch (EX Series)**—In Junos OS Release 21.1R1, we introduce the EX4400 switch, which provides connectivity for high-density environments and scalability for growing networks. The switch is available in the following models: EX4400-24T, EX4400-24P, EX4400-48T, EX4400-48P, and EX4400-48F.

  EX4400 switches support both manual and auto-channelization, but manual CLI channelization always takes precedence (see Port Settings).

  To install the EX4400 switch hardware and perform initial software configuration, routine maintenance, and troubleshooting, see EX4400 Switch Hardware Guide. See Feature Explorer for the complete list of features for any platform.

  Table 2: Feature Support on the EX4400

Feature	Description
Class of service	Support for CoS configuration with the following limitations: • If you apply strict-high priority schedulers to queues 0 through 3, then the strict-high priority schedulers are also applied to queues 8 through 11. Therefore, we recommend that you apply strict-high priority schedulers only to queues 4 through 7. • The EX4400 doesn't support the `excess-rate` configuration for schedulers. [See schedulers (CoS).]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
EVPN	Support for Layer 2 VXLAN gateway services in an EVPN-VXLAN network: • 802.1X authentication, accounting, CWA authentication, and captive portal • CoS • DHCPv4 and DHCPv6 snooping, dynamc ARP inspection (DAI), neighbor discovery inspection, IP source guard and IPv6 source guard, and router advertisement (RA) guard (no multihoming) • Firewall filters and policing • Storm control, port mirroring, and MAC filtering [See EVPN Feature Guide.]
	Support for the following Layer 2 VXLAN gateway features in an EVPN-VXLAN network: • Active/active multihoming • Proxy ARP use and ARP suppression, and Neighbor Discovery Protocol (NDP) use and NDP suppression on non-IRB interfaces • Ingress node replication for broadcast, unknown unicast, and multicast (BUM) traffic forwarding [See EVPN Feature Guide.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
	Layer 3 VXLAN gateway in EVPN-VXLAN centrally routed bridging overlay or edge-routed bridging overlay networks, supported on standalone switches or Virtual Chassis and including the following features: • Default gateway using IRB interfaces to route traffic between VLANs. [See Using a Default Layer 3 Gateway to Route Traffic in an EVPN-VXLAN Overlay Network.] • IPv6 data traffic routed through an EVPN-VXLAN overlay network with an IPv4 underlay. [See Routing IPv6 Data Traffic through an EVPN-VXLAN Network with an IPv4 Underlay.] • EVPN pure Type 5 routes. [See Understanding EVPN Pure Type-5 Routes.] The Virtual Chassis doesn't support EVPN-VXLAN multihoming, but you can use the standalone switch as an EVPN-VXLAN provider edge device in multihoming use cases.
	Support for VXLAN Group Based Policy (VXLAN-GBP). EX4400 switches support the use of existing Layer 3 VXLAN network identifiers (VNI) in conjunction with firewall filter policies to provide microsegmentation at the device or tag level, independent of the underlying network topology. IoT devices, for example, typically only need access to specific applications on the network. GBP keeps this traffic isolated by automatically applying security policies without the need for L2 or L3 lookups, or access control lists (ACLs). [See Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN.]
High availability (HA) and resiliency	High availability includes NSSU, GRES, NSB, and NSR. [See High Availability User Guide.]
Interfaces and chassis	EX4400-24T and EX4400-24P models have 24 RJ-45 ports and 2 QSFP28 ports. EX4400-48T and EX4400-48P models have 48 RJ-45 ports and 2 QSFP28 ports. The EX4400-48F model has 36 1GbE SFP ports, 12 10GbE SFP+ ports, and 2 100GbE QSFP28 ports. You can channelize the QSFP28 ports into four 25-Gbps or four 10-Gbps interfaces. [See Port Settings.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
	Support for the IEEE 802.3bt standard for Power over Ethernet (PoE) and fast PoE. With fast PoE enabled, the switch saves PoE power settings across a reboot and powers on the powered device (PD) at the initial stage of the boot (within a few seconds of switching on power) before the complete switch is booted. To configure fast PoE, use the command `set poe fast-poe`. [See Understanding PoE on EX Series Switches.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
Junos telemetry interface (JTI)	JTI Packet Forwarding Engine and Routing Engine sensor support. Use the Junos telemetry interface (JTI) and remote procedure calls (gRPC) to stream statistics from the switches to an outside collector. The following Routing Engine statistics are supported: • LACP state export • Chassis environmentals export • Network discovery chassis and components • LLDP export and LLDP model • BGP peer information (RPD) • RPD task memory utilization export • Network discovery ARP table state • Network discovery NDP table state The following Packet Forwarding Engine statistics are supported: • Congestion and latency monitoring • Logical interface • Filter • Physical interface • 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. [See Configuring a Junos Telemetry Interface Sensor (CLI Procedure), Configure a NETCONF Proxy Telemetry Sensor in Junos, and Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface).]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
Junos XML API and scripting	Support for Python, SLAX, and XSLT scripting languages and for commit scripts and macros, event policy and event scripts, op scripts, and SNMP scripts. [See Automation Scripting User Guide.]
Layer 2 features	Support for Ethernet ring protection switching version 2 (ERPSv2), which reliably achieves carrier-class network requirements for Ethernet topologies to form a closed loop. [See Example: Configuring Ethernet Ring Protection Switching on QFX Series and EX Series Switches Supporting ELS.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
Layer 2 unicast features	• Bridge protocol data unit (BPDU) protection • Ethernet ring protection switching (ERPS) • IEEE 802.1p • LAG resilient hashing • Layer 3 VLAN-tagged subinterfaces • LLDP (IEEE 802.1AB) • Loop protection • MAC address aging • MAC address filtering • Disable MAC learning • Multiple Spanning Tree Protocol (MSTP) (IEEE 802.1s) • Multiple VLAN Registration Protocol (MVRP) (IEEE 802.1ak) • Persistent MAC (sticky MAC) • Per VLAN MAC learning (limit) • Port-based VLAN • Proxy ARP • Redundant trunk group (RTG) • Root protection • Routed VLAN interface (RVI) • Rapid Spanning Tree Protocol (RSTP) (IEEE 802.1w) • Static and dynamic link aggregation with LACP (fast and slow LACP) • Static MAC address assignment for interface

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
	• Storm control • STP (IEEE 802.1D) • Uplink failure detection • VLAN • VLAN—IEEE 802.1Q VLAN trunking • VSTP [See Ethernet Switching User Guide, Security Services Administration Guide, and Spanning-Tree Protocols User Guide.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
Layer 3 unicast features	• 32-way equal-cost multipath (ECMP) • BFD (for RIP, OSPF, IS-IS, BGP, and PIM) • BGP 4-byte ASN support • BGP Add Path (BGP-AP) • Filter based forwarding (FBF) • IP directed broadcast traffic forwarding • IPv4 BGP • IPv4 multiprotocol BGP (MBGP) • IPv4 over GRE • IPv6 BGP • IPv6 CoS (BA, classification and rewrite, scheduling based on traffic class) • IPv6 IS-IS • IPv6 Neighbor Discovery Protocol (NDP) • IPv6 OSPFv3 • IPv6 ping • IPv6 stateless auto-configuration • IPv6 static routing • IPv6 traceroute • IS-IS • OSPFv2 • Path MTU discovery • RIPv2

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
	• Static routing • Unicast reverse path forwarding (unicast RPF) • Virtual router for IS-IS, RIP, OSPF, and BGP • Virtual Router Redundancy Protocol (VRRP) • VRRPv3 [See High Availability User Guide, BGP User Guide, Routing Policies, Firewall Filters, and Traffic Policers User Guide, IS-IS User Guide, Security Services Administration Guide, and OSPF User Guide.]
Licensing	You need a license to use the software features on the EX4400-24T, EX4400-24P, EX4400-48T, EX4400-48P, and EX4400-48F switches. To learn about the features supported on this device. [See EX Series Switches Support for the Juniper Flex Program.] [To add, delete, and manage licenses, see Managing Licenses.]
Multicast	• IGMP snooping • IGMP: version 1, version 2, version 3 • Multicast Listener Discovery (MLD) snooping • PIM-SM, PIM-SSM, PIM-DM [See Multicast Protocols User Guide.]
Network management and monitoring	Chef support for EX4400-48F. [See Chef for Junos OS Getting Started Guide.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
	EX4400 switches support the following Ethernet OAM link fault management (LFM) and connectivity fault management (CFM) features: • Monitor faults, using the continuity check messages (CCM) protocol to discover and maintain adjacencies at the VLAN or link level. • Discover paths and verify faults, using the Link Trace Message protocol (LTM protocol) to map the path taken to a destination MAC address. • Isolate faults, using loopback messages The EX4400 supports the following Ethernet switching events: • `adjacency loss` • `connection-protection-tlv` • `interface-status-tlv` • `port-status-tlv` EX Series switches support the `interface-down` action. [See Ethernet OAM and CFM for Switches and OAM Link Fault Management.]
	• Local and remote port mirroring, and remote port mirroring to an IP address (GRE encapsulation). [See Port Mirroring and Analyzers.] • sFlow network monitoring technology. [See sFlow Monitoring Technology.]
	Support for Puppet for Junos OS. [See Puppet for Junos OS Administration Guide.]
	Support for adding nonnative YANG modules to the Junos OS schema. [See Understanding the Management of Nonnative YANG Modules on Devices Running Junos OS.]
	Support for configuring the ephemeral database using the NETCONF and Junos XML protocols. [See Understanding the Ephemeral Configuration Database.]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
	Support for Juniper Mist Wired Assurance. You can automatically onboard and provision Juniper Networks EX4400 switches to the Juniper Mist cloud using a single activation code. Juniper Mist Wired Assurance provides automated operations and enables the use of service-level expectations (SLEs) for IoT devices, Juniper access points driven by Mist AI, and other network devices. [For an overview of Juniper Mist Wired Assurance and deployment instructions, see Juniper AI-Driven Enterprise and Overview of EX Series Switches and the Juniper Mist Cloud.]
Routing policy and firewall filters	Firewall filters and policers. [See Firewall Filters Overview.]
Security	Support for distributed denial-of-service (DDoS) protection. [See Control Plane Distributed Denial-of-Service (DDoS) Protection Overview.]
	Support for the following port security features: • DHCP snooping (IPv4 and IPv6) • Dynamic ARP inspection (DAI) • IPv6 neighbor discovery inspection [See Security Services Administration Guide.]
	Support for Media Access Control security with 256-bit cipher suite. [See Understanding Media Access Control Security (MACsec).]
Services applications	Flow-based telemetry (FBT) enables per-flow-level analytics, using inline monitoring services to create flows and collect them. A flow is a sequence of packets that have the same source IP, destination IP, source port, destination port, or protocol on an interface. For each flow, various parameters are collected and sent to a collector using the open-standard IPFIX template to organize the flow. You configure FBT by configuring the `template` statement at the `[edit services inline-monitoring]` hierarchy level, and including the `flow-monitoring` option. [See Inline Monitoring Services Configuration and template (Inline Monitoring).]

Table 2: Feature Support on the EX4400 *(Continued)*

Feature	Description
Software installation and upgrade	Support for secure boot. The implementation is based on the UEFI 2.4 standard. [See Software Installation and Upgrade Guide.]
Virtual Chassis	Virtual Chassis support for up to ten EX4400 switches interconnected and managed as a single device. The Virtual Chassis also supports NSSU to upgrade all member devices with a single command. You configure and operate an EX4400 Virtual Chassis the same way as you do other EX Series and QFX Series Virtual Chassis. However, there are a few platform-specific VCP differences, including the following: • By default, the two rear-panel 100GbE QSFP28 ports operate as four logical 50-Gbps VCP interfaces to connect the member switches. You can't use any other ports as VCPs. • These ports are in PIC slot 1, so the VCP ports on a switch are always named vcp-255/1/x, where x is a port number from 0 through 3. [See Virtual Chassis Overview for Switches.]

Authentication and Access Control

- **FQDN support in RADIUS configuration (EX2300, EX3400, EX4300, and EX4300-48P switches)**—Starting in Junos OS Release 21.1R1, RADIUS server configuration supports fully qualified domain names (FQDNs) that resolve to one or more IP addresses. This feature can be used in a cloud-managed architecture where the server name could translate to more than one IP address. RADIUS requests can be distributed across multiple servers without explicitly configuring each server IP address. Load distribution can be achieved by configuring the round-robin algorithm at the [`edit access profile profile-name radius options`] hierarchy level.

  [See Specifying RADIUS Server Connections on Switches.]

EVPN

- **Tunnel endpoint in the PMSI tunnel attribute field for EVPN Type 3 routes (ACX5448, EX4600, EX4650, EX9200, and QFX10002)**—Starting in Junos OS Release 21.1R1, you can set the tunnel endpoint in the provider multicast service interface (PMSI) tunnel attribute field to use the ingress router's secondary loopback address. When you configure multiple loopback IP addresses on the local provider edge (PE) router and the primary router ID is not part of the MPLS network, the remote PE router cannot set up a PMSI tunnel route back to the ingress router.

To configure the router to use a secondary IP address that is part of the MPLS network, include the `pmsi-tunnel-endpoint`*`pmsi-tunnel-endpoint`* statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn]` hierarchy level for both EVPN and virtual-switch instance types.

[See EVPN.]

- **Flow-aware transport pseudowire support for EVPN-VPWS (MX Series routers and EX9200 switches)**—Starting in Junos OS Release 21.1R1, you can statically configure provider edge (PE) devices to use flow-aware transport (FAT) pseudowire labels in an EVPN virtual private wire service (VPWS) routing instance with an IP/MPLS underlay fabric. PE devices use these labels to load-balance EVPN-MPLS packets across ECMP paths or link aggregation groups (LAGs) without needing to do deep packet inspection of the payload.

  To enable FAT pseudowire load balancing in an `evpn-vpws` routing instance:

  - Configure `flow-label-transmit-static` on PE devices to insert FAT flow labels into VPWS pseudowire packets sent to remote PE devices.
  - Configure `flow-label-receive-static` on PE devices to remove FAT flow labels from VPWS pseudowire packets received from remote PE devices.

  You can configure these statements for all pseudowires in the routing instance or for pseudowires associated with a specific interface in the routing instance.

  [See FAT Flow Labels in EVPN-VPWS Routing Instances, flow-label-receive-static, and flow-label-transmit-static.]

- **EVPN-VXLAN fabric (EX9200 switches with EX9200-15C line cards)**—Starting in Junos OS Release 21.1R1, EX9200 switches with EX9200-15C line cards support the following features in an EVPN-VXLAN fabric:

  - Layer 2 VXLAN
    - Multihoming in active/active mode, an Ethernet segment identifier (ESI) per interface, and preference-based designated forwarder (DF) election
    - MAC pinning, MAC move, MAC limiting, and MAC aging
    - QoS
    - DHCP and DHCP relay
    - Prevention of broadcast, unknown unicast, and multicast (BUM) traffic loops when a leaf device is multihomed to more than one spine device
  - Layer 3 VXLAN
    - IRB interfaces

    - IPv6 over IRB interfaces
    - Support for OSPF, IS-IS, BGP, and static routing over IRB interfaces
    - Proxy ARP and ARP suppression, and proxy NDP and NDP suppression with and without IRB interfaces
    - IPv6 underlay
    - Virtual machine traffic optimization (VMTO) for ingress traffic
  - Data Center Interconnect (DCI)
    - Pure EVPN Type 5 routes only
  - High availability
    - Nonstop active routing (NSR)
    - GRES
    - Graceful restart from a routing process restart or Routing Engine switchover without NSR enabled
  - Operations and management
    - Core isolation feature
    - Ping over EVPN Type 5 tunnel
  - Static VXLAN
    - Overlay ping and traceroute

  [See EVPN User Guide.]

- **Loop detection for EVPN-VXLAN fabrics (EX4300-48MP)**—Starting in Junos OS Release 21.1R1, you can configure loop detection on the server-facing Layer 2 interfaces on EX4300-48MP leaf devices in an EVPN-VXLAN fabric. This feature can detect the following types of Ethernet loops:
  - A loop between two interfaces with different Ethernet segment identifiers (ESIs), usually caused if you miswire fabric components.
  - A loop between two interfaces with the same ESI, usually caused if you miswire a third-party switch to the fabric.

  After you enable loop detection, the interfaces periodically send multicast loop-detection protocol data units (PDUs). If a loop detection-enabled interface receives a PDU, the device detects a loop, which triggers the configured action to break the loop. For example, if you configure the `interface-down`

action, the device brings down the interface. After the `revert-interval` timer expires, the device reverts the action and brings the interface back up again.

[See loop-detect.]

- Macro Segmentation using Group Based Policy (EX4400)—Starting in Junos OS Release 21.1R1. Juniper EX4400 series switches support the use of existing layer 3 VXLAN network identifiers (VNI) in conjunction with firewall filter policies to provide micro-segmentation at the level of device or tag, independent of the underlying network topology. IoT devices, for example, typically only need access to specific applications on the network. VXLAN-GBP can keep this traffic isolated by automatically applying security policies without the need for L2 or L3 lookups or ACLs.

  To use VXLAN-GBP, enable group-based policies at the global hierarchy level: `[chassis forwarding-options vxlan-gbp-profile]` on the tunnel termination endpoint. Two new match condition terms, `gbp-src-tag` and `gbp-dst-tag`, are introduced at the `[firewall family ethernet-switching filter` *`name`* `term` *`name`* `from ]` level of the hierarchy. You can configure the tags manually or get them from the RADIUS server upon authentication by the host.

  [See https://www.juniper.net/documentation/us/en/software/junos/evpn-vxlan/topics/example/micro-segmentation-using-group-based-policy.html Example: Micro and Macro Segmentation using Group Based Policy in a VXLAN.]

- **Explicit congestion notification (ECN) over VXLAN tunnels (EX4650 and QFX5120)**—Starting in Junos OS Release 21.1R1, by default, standalone EX4650 and QFX5120 switches support explicit congestion notification (ECN) for packets that are encapsulated across VXLAN tunnels, as follows:

  - During VXLAN encapsulation at the source virtual tunnel endpoint (VTEP), the switch copies the ECN bits of the Type-of-Service (ToS) field from the original packet IP header to the outer VXLAN encapsulation IP header.

  - During VXLAN de-encapsulation at the remote VTEP, the switch copies the ECN bits of the ToS field from the outer VXLAN encapsulation IP header to the original packet IP header.

  You can configure the `vxlan-disable-copy-tos-encap` statement or the `vxlan-disable-copy-tos-decap` statement at the `[edit forwarding-options]` hierarchy on the encapsulation or de-encapsulation ends of the tunnel, respectively, to disable the ECN copy operation.

  > **NOTE**: These switches also copy the differentiated services code point (DSCP) bits in the ToS field of the IP header upon VXLAN encapsulation and de-encapsulation by default, and the same statements disable copying both the DSCP and ECN bits.

  [See vxlan-disable-copy-tos-encap and vxlan-disable-copy-tos-decap.]

Forwarding Options

- **Storm control support in EVPN-VXLAN overlay networks (EX4650 switches)**—Starting in Junos OS Release 21.1R1, EX4650 switches support storm control in an EVPN-VXLAN overlay network. Storm control enables the switch to monitor traffic levels and to drop broadcast, unknown unicast, and multicast (BUM) packets before they cause a traffic storm.

  [See Understanding Storm Control.]

High Availability

- **Support for VRRP on EX9200-SF3 and EX9200-15C (EX9200)**—Starting in Junos OS Release 21.1R1, the EX9200-SF3 Switch Fabric module and the EX9200-15C line card support VRRP. All VRRP features are supported.

  [See Understanding VRRP.]

Licensing

- **Juniper Agile Licensing (EX2300, EX2300-MP, EX2300-C, EX3400, EX4300, EX4300-MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, and EX4400-48T)**—Starting in Junos OS Release 21.1R1, the listed EX Series switches support Juniper Agile Licensing.
  Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can use Juniper Agile Licensing to install and manage licenses for hardware and software features.

  Juniper Agile Licensing supports soft enforcement and hard enforcement of hardware and software feature licenses.

  - With soft enforcement, if you configure a feature without a license, Junos OS displays a warning when you commit the configuration. However, the feature remains operational. In addition, Junos OS generates periodic alarms indicating that you need the license to use the feature. You can see the list of alarms at System Log Explorer.

  - With hard enforcement, if you configure a feature without a license, Junos OS displays a warning when you commit the configuration. The feature is not operational until the license is installed. In addition, Junos OS generates periodic alarms indicating that you need the license to use the feature. You can see the list of alarms at System Log Explorer.

  "Table 3" on page 45 describes the licensing support for soft-enforced features on EX2300 switches.

Table 3: Licensed Features on EX2300 switches

License Model	Use Case Examples or Solutions	Feature List
Standard	Campus and access Layer 2 or Layer 3	• Layer 2 and Layer 3 filters • Layer 2 (*x*STP, 802.1Q, and LAG) • Layer 2 and Layer 3 QoS • Layer 3 (static) • IGMP snooping • Operation, Administration, and Maintenance (OAM) link fault management (LFM) • sFlow • SNMP • Junos telemetry interface (JTI) • Virtual Chassis*

Table 3: Licensed Features on EX2300 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List
Advanced	Campus and access Layer 2 or Layer 3	• Bidirectional Forwarding Detection (BFD) • IGMP version 1, IGMP version 2, and IGMP version 3 • IPv6 routing protocols: Multicast Listener Discovery (MLD) version 1 and MLD version 2, OSPF version 3, PIM multicast, VRRP version 3 • Multicast Source Discovery protocol (MSDP) • OAM and Maintenance CFM • OSPF version 2 or OSPF version 3 • Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode • Real-time performance monitoring (RPM) • RIP IPv6 (RIPng) • VRRP

Virtual Chassis*—We've included Virtual Chassis license in the Standard license model on EX2300-C 12-port switches. However, we don't include the Virtual Chassis license on EX2300 24-port and 48-port switch models. You need to purchase the license separately.

"Table 4" on page 47 describes the licensing support for soft-enforced features on EX3400 switches.

Table 4: Licensed Features on EX3400 switches

License Model	Use Case Examples or Solutions	Feature List
Standard	Campus and access Layer 2 or Layer 3	• Layer 2 and Layer 3 filters • Layer 2 (*x*STP, 802.1Q, and LAG) • Layer 2 and Layer 3 QoS • Layer 3 (static) • IGMP snooping • Operations, Administration, and Maintenance (OAM) link fault management (LFM) • sFlow • SNMP • Junos telemetry interface (JTI) • Virtual Chassis

Table 4: Licensed Features on EX3400 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List
Advanced	Campus and access Layer 2 or Layer 3	• Bidirectional Forwarding Detection (BFD) • IGMP version 1, IGMP version 2, and IGMP version 3 • IPv6 routing protocols: Multicast Listener Discovery (MLD) version 1 and MLD version 2, OSPF version 3, PIM multicast, VRRP version 3 • Multicast Source Discovery protocol (MSDP) • OAM CFM • OSPF version 2 or OSPF version 3 • Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode • Real-time performance monitoring (RPM) • RIP IPv6 (RIPng) • Unicast reverse-path forwarding (unicast RPF) • Virtual router • VRRP

Table 4: Licensed Features on EX3400 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List
Premium	Campus and access Layer 3	• Bidirectional Forwarding Detection (BFD) • IGMP version 1, IGMP version 2, and IGMP version 3 • IPv6 routing protocols: Multicast Listener Discovery (MLD) version 1 and MLD version 2, OSPF version 3, PIM multicast, VRRPv3, virtual router support for unicast and filter-based forwarding (FBF) • Multicast Source Discovery Protocol (MSDP) • OAM CFM • OSPF version 2 or OSPF version 3 • Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode • Real-time performance monitoring (RPM) • RIP IPv6 (RIPng) • Unicast reverse-path forwarding (unicast RPF) • Virtual router • VRRP • BGP and multiprotocol BGP (MBGP) • IS-IS

"Table 5" on page 50 describes the licensing support for soft-enforced features on EX4300 switches.

Table 5: Licensed Features on EX4300 switches

License Model	Use Case Examples or Solutions	Feature List
Standard	Campus and access Layer 2 or Layer 3	• Layer 2 and Layer 3 filters • Layer 2 (*x*STP, 802.1Q, and LAG) • Layer 2 and Layer 3 QoS • Layer 3 (static) • IGMP snooping • Operations, Administration, and Maintenance (OAM) link fault management (LFM) • sFlow • SNMP • Junos telemetry interface (JTI) • Virtual Chassis

Table 5: Licensed Features on EX4300 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List
Advanced	Campus and access Layer 2 or Layer 3	• Bidirectional Forwarding Detection (BFD) • IGMP version 1, IGMP version 2, and IGMP version 3 • Multicast Source Discovery protocol (MSDP) • OAM CFM • OSPF version 2 or OSPF version 3 • Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode • Real-time performance monitoring (RPM) • RIP IPv6 (RIPng) • Unicast reverse-path forwarding (unicast RPF) • Virtual router • VRRP

Table 5: Licensed Features on EX4300 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List

Premium	Campus and access Layer 3	• Bidirectional Forwarding Detection (BFD)
		• CFM (IEEE 802.1ag)
		• IGMP version 1, IGMP version 2, and IGMP version 3
		• Multicast Source Discovery Protocol (MSDP)
		• OAM CFM
		• OSPF version 2 or OSPF version 3
		• Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode
		• Real-time performance monitoring (RPM)
		• RIP IPv6 (RIPng)
		• Unicast reverse-path forwarding (unicast RPF)
		• Virtual router
		• VRRP
		• BGP and multiprotocol BGP (MBGP)
		• IS-IS
		• EVPN-VXLAN
		  • Supported only on EX4300-48MP switch.
		  • Requires the BGP for configuration.

"Table 6" on page 53 describes the licensing support for soft-enforced features on EX4400 switches.

Table 6: Licensed Features on EX4400 switches

License Model	Use Case Examples or Solutions	Feature List
Standard	Campus and access Layer 2 or Layer 3	• Layer 2 and Layer 3 filters • Layer 2 (*x*STP, 802.1Q, and LAG) • Layer 2 and Layer 3 QoS • Layer 3 (static) • IGMP snooping • Operations, Administration, and Maintenance (OAM) link fault management (LFM) • sFlow • SNMP • Junos telemetry interface (JTI) • Virtual Chassis

Table 6: Licensed Features on EX4400 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List
Advanced	Campus and access Layer 2 or Layer 3	• Bidirectional Forwarding Detection (BFD) • IGMP version 1, IGMP version 2, and IGMP version 3 • Multicast Source Discovery protocol (MSDP) • OAM CFM • OSPF version 2 or OSPF version 3 • Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode • Real-time performance monitoring (RPM) • RIP IPv6 (RIPng) • Unicast reverse-path forwarding (unicast RPF) • Virtual router • VRRP

Table 6: Licensed Features on EX4400 switches *(Continued)*

License Model	Use Case Examples or Solutions	Feature List
Premium	Campus and access Layer 3	• Bidirectional Forwarding Detection (BFD) • CFM (IEEE 802.1ag) • IGMP version 1, IGMP version 2, and IGMP version 3 • Multicast Source Discovery Protocol (MSDP) • OAM CFM • OSPF version 2 or OSPF version 3 • Protocol Independent Multicast (PIM) dense mode, PIM source-specific mode, PIM sparse mode • Real-time performance monitoring (RPM) • RIP IPv6 (RIPng) • Unicast reverse-path forwarding (unicast RPF) • Virtual router • VRRP • BGP and multiprotocol BGP (MBGP) • IS-IS • EVPN-VXLAN • Requires the BGP for configuration.

On EX4400 switch, the flow-based telemetry and MACsec features are hard-enforced. You'll need a license to use these features.

The flow-based telemetry and MACsec features are supported only in standalone mode.

[See Supported Features on EX2300, EX2300-MP, EX2300-C, EX3400, EX4300, EX4300-MP, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, and EX4400-48T devices, Juniper Agile Licensing Guide, and Configuring Licenses in Junos OS.]

Network Management and Monitoring

- **Ephemeral configuration database support for load update operations (EX9200, MX5, MX10, MX80, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMX)**—Starting in Junos OS Release 21.1R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database using a load update operation. To perform a load update operation, set the `<load-configuration> action` attribute to `update`.

  [See <load-configuration>.]

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

  [See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

- **Remote port mirroring to IPv6 address (GRE encapsulation)(EX4650, EX4650-48Y-VC, QFX5120, QFX5120-32C, QFX511120-48T, QFX5120-48T-VC, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 21.1R1, you can use remote port mirroring to copy packets entering a port or VLAN and sends the copies to the IPv6 address of a device running an analyzer application on a remote network (sometimes referred to as “extended port mirroring”). When you use remote port mirroring the mirrored packets are GRE-encapsulated.

  Add the address you would like to have the copied packets sent to in the CLI hierarchy. For example, `set forwarding-options analyzer ff output ipv6-address 2000::1`.

  [See Understanding Port Mirroring and Analyzers.]

Software Installation and Upgrade

- **Support for bootstrapping using HTTP proxy server in phone-home client (EX2300, EX2300-VC, EX3400, EX3400-VC, EX4400-24T, EX4400-48F, EX4400-48T, and EX4600)**—Starting in Junos OS Release 21.1R1, when the phone-home client (PHC) receives information regarding the HTTP proxy server through either DHCP option 43 suboption 8 or DHCP option 17 suboption 8, it creates an HTTPS transparent tunnel with the proxy server. After the tunnel is established, the PHC uses the

tunnel as a proxy for the phone-home server or redirect server. The phone-home client downloads the software image and configuration file through the tunnel onto the device. When bootstrapping is complete, the device reboots and the tunnel quits.

[See Obtaining Configurations and Software Image Without User Intervention Using Phone-Home Client.]

- **Support for DHCP option 43 suboption 8 to provide proxy server information in phone-home client (EX2300, EX2300-VC, EX3400, EX3400-VC, EX4400-24T, EX4400-48F, EX4400-48T, and EX4600)**—Starting in Junos OS Release 21.1R1, during the bootstrapping process, the phone-home client (PHC) can access the redirect server or the phone-home server through a proxy server. The DHCP server uses DHCP option 43 suboption 8 or DHCP option 17 suboption 8 to deliver the details of both IPv4 and IPv6 proxy servers to the PHC. The DHCP daemon running on the target switch learns about the proxy servers in the initial DHCP cycle and then populates either the phc_vendor_specific_info.xml files or the phc_v6_vendor-specific_info.xml files located in the **/var/etc/** directory with the vendor-specific information.

  [See Obtaining Configurations and Software Image Without User Intervention Using Phone-Home Client.]

- **Support for phone-home client (EX4400 Virtual Chassis)**—Starting in Junos OS Release 21.1R1, the phone-home client (PHC) can securely provision an EX4400 Virtual Chassis without requiring user interaction. You only need to:

  - Ensure that the Virtual Chassis members have the factory-default configuration.
  - Interconnect the member switches using dedicated or default-configured Virtual Chassis ports.
  - Connect the Virtual Chassis management port or any network port to the network.
  - Power on the Virtual Chassis members.

  The PHC automatically starts up on the Virtual Chassis and connects to the phone-home server (PHS). The PHS responds with bootstrapping information, including the Virtual Chassis topology, software image, and configuration. The PHC upgrades each Virtual Chassis member with the new image and applies the configuration, and the Virtual Chassis is ready to go.

  [See Provision a Virtual Chassis Using the Phone-Home Client.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for EX Series switches.

What's Changed in Release 21.1R3

IN THIS SECTION

EVPN

- **Output for show Ethernet switching flood extensive**--The output for `show ethernet-switching flood extensive` now displays the correct next-hop type for Virtual Ethernet and WAN mesh group in an EVPN-VXLAN network as unilist. Previously, the output for show ethernet-switching flood extensive would misidentify the next-hop type as composite.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair need to be unique only among the FTI tunnels of the same tunnel encapsulation type. Prior to this PR, the source and destination address pair had to be unique among all the FTI tunnels regardless of the tunnel encapsulation type.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

  [See request system scripts refresh-from and cert-file (Scripts).]

What's Changed in Release 21.1R2

IN THIS SECTION

EVPN

- **IGMP snooping options has changed hierarchy level**—Junos OS has moved the following options from the `[edit protocols igmp-snooping]` hierarchy to `[edit protocols igmp-snooping vlan vlan-name]` hierarchy and `[edit routing-instances evpn protocols igmp-snooping]` hierarchy to `[edit routing-instances evpn protocols igmp-snooping vlan vlan-name]` hierarchy

- query-interval

  query-last-member-interval

  query-response-interval

  robust-count

evpn-ssm-reports-only

immediate-leave

- **Support for displaying SVLBNH information**—You can now view shared VXLAN load balancing next hop (SVLBNH) information when you display the VXLAN tunnel endpoint information for a specified ESI and routing instance by using `show ethernet-switching vxlan-tunnel-end-point esi esi-identifier esi-identifier instance instance svlbnh` command.

General Routing

- **Configure internal IPsec authentication algorithm (EX Series)**—You can configure the algorithm `hmac-sha-256-128` at the `[edit security ipsec internal security-association manual direction bidirectional authentication algorithm]` hierarchy level for internal IP security (IPsec) authentication. In earlier releases, you could configure the algorithm `hmac-sha-256-128` for MX Series devices only.

Interfaces and Chassis

- **Blocking duplicate IP detection in the same routing instance (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**--Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. Refer to the table mentioned in the topic inet (interfaces). When you try to configure same IP on two logical interfaces inside same routing instance, the commit will be blocked with the error displayed as shown below:[edit] `user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24` [edit] `user@host# commit commit complete` [edit] `user@host# set interfaces ge-0/0/2 unit 0 family inet address 2.2.2.2/24` [edit] `user@host# commit` [edit interfaces ge-0/0/2 unit 0 family inet] 'address 2.2.2.2/24' identical local address found on rt_inst **default**, intfs **ge-0/0/2.0 and ge-0/0/1.0**, family **inet**. error: configuration check-out failed.

  [See inet(interfaces).]

Layer 2 Ethernet Services

- **Link selection support for DHCP**—We have introduced the `link-selection` statement at the **`edit forwarding-options dhcp-relay relay-option-82`** hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Prior to this release, the DHCP relay dropped packets during the renewal DHCP process and the DHCP server used the leaf's address as a destination to acknowledge the DHCP renewal message.

  See relay-option-82.

Network Management and Monitoring

- **Change in OID ifHighSpeed**—Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

  See SNMP MIBs and Traps Supported by Junos OS.

- **Changes in contextEngineID for SNMPv3 INFORMS (PTX Series, QFX Series, ACX Series, EX Series, MX Series, and SRX Series**— Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  See SNMP MIBs and Traps Supported by Junos OS.

What's Changed in Release 21.1R1

IN THIS SECTION

EVPN

- **IGMP snooping options has changed hierarchy level--**Junos OS has moved the following options from the `edit protocols igmp-snooping` hierarchy to `edit protocols igmp-snooping vlan vlan-name vlan-all` hierarchy and `edit routing-instances evpn protocols igmp-snooping` hierarchy. to `edit routing-instances evpn protocols igmp-snooping vlan vlan-name vlan-all` hierarchy

  - query-interval

    query-last-member-interval

    query-response-interval

    robust-count

evpn-ssm-reports-only

immediate-leave

General Routing

- **Change in license bandwidth command on vMX virtual routers**—Starting in Junos OS, to use the available license bandwidth, explicitly set the license bandwidth using the `set chassis license bandwidth in mbps` command.

  [See Configuring Licenses on vMX Virtual Routers.]

- **Configure internal IPsec authentication algorithm (EX Series)**—You can configure the algorithm `hmac-sha-256-128` at the `edit security ipsec internal security-association manual direction bidirectional authentication algorithm` hierarchy level for internal IP security (IPsec) authentication. In earlier releases, you could configure the algorithm `hmac-sha-256-128` for MX Series devices only.

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX event scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Layer 2 Ethernet Services

- **Modification to sync-reset command (All JUNOS and EVO platforms)**—Starting from this release, the sync-reset command is disabled by default on all Junos and EVO platforms. Sync-reset command enables the device to send the sync bit in the LACP packets on minimum-link failure. Previously the sync-reset command was enabled by default on QFX and EX series, while it was by default disabled on MX, PTX and ACX series.

  [See sync-reset.]

Network Management and Monitoring

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the `[edit system services netconf hello-message yang-module-capabilities]` hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the `[edit system services netconf netconf-monitoring netconf-state-schemas]` hierarchy level.

  [See hello-message and netconf-monitoring.]

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the `verbose` statement at the `[edit system export-format json]` hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from `verbose` to `ietf` starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the

appropriate statement at the `[edit system export-format json]` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See export-format.]

Known Limitations

Learn about known limitations in this release 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.

EVPN

- After a reboot during recovery process, the ESI LAGs come up before the BGP sessions, and routes or ARP entries are not synchronized. PR1487112

General Routing

- When the device is up and running for a long time, there is a possibility FS gets bad blocks and it is accumulated. When any change done to it, it reloads and tries to recover the bad blocks from the FS. PR910445
- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. PR1385970
- On all platforms running Junos OS, in a Q-in-Q environment, if xSTP is enabled on an interface that has a logical interface with vlan-id-list configured, then, it will only run on those logical interfaces whose vlan-id range includes native-vlan-id configured. All other xSTP will be in discarding state. This might lead to traffic drop. PR1532992
- Whenever a firewall is configured with action inline monitoring and to count, the corresponding firewall counters will not show any increment in traffic statistics. This limitation is from underlying hardware. PR1542192
- Inline monitoring security service might not report a IPv4 security violation when traffic is received with source IP address same as destination IP address. This is limitation from Hardware. PR1542213
- On a virtual chassis, the CLI command to add license is not available on the backup member. Licenses must be added from the master member only. PR1545075

- When user issues `request support information` to collect debug data for sharing with tech support, commands which are not supported for the platform will display a syntax error message. PR1547835
- In a Virtual Chassis environment, whenever there is a Link Aggregation Group (LAG) configured with child members from different virtual chassis members, then layer2-header hashing should be used to distribute traffic across the LAG members. Usage of vlan-id based hashing is not recommended and might result in reduced throughput of the LAG. PR1548859
- Tail drop is seen in WRED configuration statistics. PR1549910
- Inline-monitoring service supports only upto a maximum of 8 instances. PR1550014
- VRRP delegate-processing currently not supported on EX4400 series of switches. PR1552076
- Packets mirrored through analyzer with ingress and egress interfaces across different virtual chassis members might have a different vlan-id from that of vlan-id of the exiting egress interface. This limitation is from underling hardware. PR1552905
- "Resource deadlock avoided" messages are observed during software add-on EX4400. No functionality impact seen. PR1557468
- License keys should be installed using either operational command or configuration set commands. Same license key should not be added via both operational and set comands. PR1557980

Interfaces and Chassis

- **Support for low power idle mode (EX4400-48T, EX4400-48P, EX4400-24T, and EX4400-24P)**— Starting in Junos OS Release 21.1R1, the 1-Gbps or 100-Mpbs port switches to low power idle (LPI) mode based on the following conditions:.
  - When a port operates at 1-Gbps speed and no traffic is either received or transmitted, then the port enters LPI mode. If the 1-Gbps port transfers unidirectional or bidirectional traffic, then the port will not enter LPI mode.
  - When a port operates at 100-Mbps speed, the port switches to LPI mode, based on the direction of the traffic. The `show interfaces interface-name extensive` command displays `RX LPI` when there is no RX traffic and `TX LPI` when there is no TX traffic.

  You can view the interface that is in LPI mode by executing the `show interfaces interface-name extensive` command. The output field `IEEE 802.3az Energy Efficient Ethernet` displays the status of the LPI mode.

  [See show interfaces extensive .]

Infrastructure

- Software versions 21.1 and lower are running FreeBSD version 11 whereas from version 21.2 onward, the FreeBSD version is 12. Software upgrade to 21.2 (or later) from 21.1 (or earlier) will mandatorily need cli knob 'no-validate' to be used during software image upgrade process. (For EX4400 platforms, this is applicable from version 21.3 onward. Hence, for EX4400 platforms, software upgrade to 21.3 (or later) from 21.2 (or earlier) will mandatorily need cli knob 'no-validate' to be used during software image upgrade process.) PR1586481

User Interface and Configuration

- Unsupported options can be seen under "restart" command. PR1545558
- Python script is not supported in ZTP workflow. Python can run (during ZTP) only in few EX Series based flex images. PR1547557

Open Issues

IN THIS SECTION

Learn about open issues in this release 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.

EVPN

- On all Junos OS platforms with EVPN scenario, the number of MAC-IP binding counters might reach the limit when MAC-IP is moved between interfaces. Since MAC-IP counters are not decremented when entry is deleted due to this defect, repeated moves will result in a limit (default value is 1024) that will be reached even though there are fewer entries. Meanwhile, traffic loss might be seen. PR1591264

Forwarding and Sampling

- `fast-lookup-filter` with match not supported in FLT hardware might cause the traffic drop. PR1573350

General Routing

- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. PR1362609
- On EX9208 switch, a few xe- interfaces go down with an error message "if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC error". PR1377840
- On EX9214 device, the following error message is observed after rebooting and MACsec-enabled link flaps: "errorlib_set_error_log(): err_id(-1718026239)". PR1448368
- When running the command `show pfe filter hw filter-name` *`filter-name`*, the command fails to retrieve the Packet Forwarding Engine programming details of the filter. PR1495712
- When a VLAN member is specified as a string, the 'IF_MSG_IFL_VADDR' TLV is not generated with the VLAN infoformation, and the TRIO afttriostream is not updated with the nativevlanId and nativevlanenable flags. Thus, the packet is treated as untagged, and when it reaches the trunk egress interface, it is dropped because the trunk interface does not allow untagged traffic to pass through. The issue is specific to platforms with ZT line cards, including EX9200-SF3 and EX9200-15C. PR1506403
- A delay of 35 seconds is added in reboot time in Junos OS Release 20.4R3 compared to Junos OS Release 19.4R2. PR1514364
- License daemon will restart and start providing the required support when intermittent license-check.core file is seen during the device initialization. There is no service impact. PR1545175

- When ICMP packets egress from the device, that might take Best-Effort queue. This avoids congestion case protocol flap when huge number of ICMP traffic being generated. Other control protocol such as OSPF and others take network-control queue, which is in parity with other QFX5000 line of switches. PR1550293
- When dot1x `server-fail-voip vlan-name` is configured, ensure that both `server-fail-voip vlan-name` and `voip vlan` are configured using vlan name and not by using vlan-id. PR1561323
- Observing traffic drop during unified ISSU due to LAG interface flap. PR1569578
- BUM traffic replication over VTEP is sending out more packets than expected and there seems to be a loop also in the topology. PR1570689
- On all Junos OS platforms, traffic loss might be observed due to a rare timing issue when performing frequent interface bridge domain (IFBD) configuration modifications. This behavior is seen when the Packet Forwarding Engine receives out-of-order IFBD(s) from Routing Engine and might lead to the fxpc process crash and traffic drop. PR1572305
- Pause frames counters are not getting incremented when pause frames are sent. PR1580560
- On EX Series switches such as EX2300, EX3400, EX4300, EX4600, and EX4650 with broadcom chip as Packet Forwarding Engine, if IS-IS is enabled on an integrated routing and bridging (IRB) interface and the maximum transmission unit (MTU) size of the IRB interface is configured with a value great than 1496 bytes, the IS-IS hello (IIH) PDUs with jumbo frame size (that is, great than 1496 bytes) might be dropped and not sent to the IS-IS neighbors. The following is the product list of EX Series switches with broadcom chip as the Packet Forwarding Engine. PR1595823
- On EX4600, after performing an upgrade, the peer device is rebooted or the peer interface is disabled and then enabled. As a result, the SFP-T port on EX4600 might remain in UP state but might not forward traffic. PR1600291
- Observing pfex core file at 0x01fdf324 in pfe_bcm_ifd_mac_config **/../src/pfe/common/pfe-arch/broadcom/applications/l2/pfe_bcm_l2_intf** while cleanup of the configurations after NSSU. PR1602873
- On EX Series line of switches, the system reboot takes approximately 9 minutes for FPCs to come online after system reboot command is issued. PR1605002
- After performing ZTP, default configuration under ge-0/0/* will be missing in EX4600 product. PR1614098

Infrastructure

- On EX Series switches except EX4300, EX4600, and EX9200, an interface is configured for single vlan or multiple vlans, if all these vlans of this interface have igmp-snooping enabled, then this interface will drop Hot Standby Router Protocol for IPv6 (HSRPv2) packets. But if some vlans do not have igmp-snooping enabled, then this interface is working fine. PR1232403
- On EX Series switches, If you are configuring a large-scale number of firewall filters on some interfaces, the FPC might crash and generate core files. PR1434927
- IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151) error message is observed continuously in AD with base configurations. PR1485038
- A double free vulnerability in the software forwarding interface daemon (sfid) process of Juniper Networks Junos OS allows an adjacently-connected attacker to cause a Denial of Service (DoS) by sending a crafted ARP packet to the device. Refer to https://kb.juniper.net/JSA11162 for more information. PR1497768
- On EX4400 family of devices, sometimes login prompt is not shown after the login session ends. PR1582754
- On EX4400 device, the cli command `show system processes detail` will not display CPU details under the CPU column. PR1588150

Layer 2 Features

- On EX Series line of switches, memory leak might be seen because of the eswd daemon that displays the following system log message: eswd[1330]: JTASK_OS_MEMHIGH: Using 212353 KB of memory, 158 percent of available /kernel: KERNEL_MEMORY_CRITICAL: System low on free memory, notifying init (#2). /kernel: Process (1254,eswd) has exceeded 85% of RLIMIT_DATA: used 114700 KB Max 131072 KB. PR1262563
- On EX4600 platforms, if a change related to TPID is made in the device control daemon, a traffic drop might be seen in the Packet Forwarding Engine due to failure on Layer 2 learning or interfaces flapping. PR1477156

Platform and Infrastructure

- On EX4300 POE switches, the pfex process CPU utilization becomes high after 6-8 weeks. There is no functional impact. PR1453107

- When the dhcp relay mode is configured as `no-snoop`, the offer gets dropped due to incorrect ASIC programing. PR1530160
- On EX9200 platforms, FPC gets restarted and thereby disrupting traffic when there is an out-of-order filter state and its terms. This issue might be seen only in back-to-back GRES in more than 40 to 50 iterations. PR1579182

Routing Policy and Firewall Filters

- On all Junos OS platforms with `set policy-options rtf-prefix-list` configured, if upgrade to a specific version, the device might fail to validate its configuration which eventually causing rpd to crash unexpectedly due to a software fault. PR1538172

User Interface and Configuration

- The issue is seen on EX Series VC only which can be avoided with a simple workaround as to providing a valid package during upgrade command. PR1557628

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases 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: 21.1R3

IN THIS SECTION

EVPN

- Traffic loss might be seen under EVPN scenario when MAC-IP moves from one CE interface to another. PR1591264
- On all Junos OS platforms, a traffic loss might be seen if an aggregated Ethernet bundle interface with ESI is disabled on a primary Routing Engine followed by a Routing Engine switchover. PR1597300

General Routing

- The power over Ethernet (POE) might not be detected and fail to work on Virtual Chassis (VC) members. This occurs because of a rare timing issue in the Virtual Chassis scenario.. PR1539933
- The Virtual Chassis Port (VCP) might not come up EX4600 platform.PR1555741
- On EX Series VC platforms, the new primary Routing Engine post switchover might go into DB mode (or crash). PR1565213
- On EX Series line of switches, in EVPN-VxLAN scenario, when there is a MAC move from vtep to local, the MAC address will not be aged out when traffic with the same src mac is stopped. PR1565624

- On all Junos OS platforms if 802.1X authentication is configured globally using the `set protocol dot1x interface all` command and if trunk interface is configured with vlans, then the private VLAN configuration might fail. PR1574480
- On EX4400, with this issue multicast queue statistics will be cumulatively reported with unicast queues. MC8-to-MC11 queues statistics will be reported under UC0 to UC3 queues. PR1578375
- On Junos OS EX Series platforms, the dcpfe crash might be seen because of the interface flap on which a large number of MAC-based VLAN clients are registered. PR1578859
- On EX Series line of swithces, some 40G ports might not be channelized successfully. PR1582105
- On EX2300 and EX3400 platforms, when using the command `request system firmware upgrade poe fpc-slot all-members` to upgrade the Power over Ethernet (PoE) firmware, the upgrade might fail. PoE firmware upgrade failure might cause PoE not to function as expected. This is a rare, intermittent issue seen very randomly. PR1584491
- On EX9204, EX9208, and EX9214 switches with EX9200-40XS linecard enabled, issuing the `set chassis fpc 1 power on` statement and commiting, results in packet drop even when the traffic is not passing over the EX9200-40XS linecard. This is a timing issue and does not occur frequently. PR1586740
- The DHCP relay might not work if it connects with the server through type 5 route which with aggregated Ethernet interface as the underlay interface. PR1592133
- On the EX4300-48MP Virtual Chassis, the backup Routing Engines clear the reporting alarm for a PEM failure intermittently for a missing power source. PR1593795
- Clients authentication failure might occur due to dot1x daemon memory leak. PR1594224
- In the EVPN/VXLAN scenario, the label field for the EVPN Type 1 route is set to 1. PR1594981
- There is a steady increase in storage usage in the backup chassis when the subscriber service is enabled. PR1595238
- The MAC/IP withdraw route might be suppressed by rpd in the EVPN-VxLAN scenario. PR1597391
- The backup Virtual Chassis member might not learn MAC address on a primary after removing the VLAN unit from the SP style aggregated Ethernet interface which is part of multiple VLAN units. PR1598346
- The l2ald process might crash because of the memory leak when all active interfaces in a VLAN are unstable. PR1599094
- On a EX4400 Virtual Chassis operating with scaled configurations and traffic, the line card console might fail to redirect to the current virtual chassis master member. PR1599625
- On the EX4300-MP switch, unable to disable the management port em1. PR1600905

- On EX2300 and EX4650, if the system is upgraded from Junos OS Release 20.2 or earlier release to Junos OS 20.3 or later release, either using phone-home feature or when the system is in factory default state, the upgrade will fail with phone-home crash. PR1601722
- On EX2300 and EX2300-MP Virtual Chassis platforms ARP might not get resolved. PR1602003
- On EX4400 dot1x authentication might not work on EVPN/xlan enabled endpoints. PR1603015
- The NSSU performed with MACsec configuration might result in generating fxpc core file. PR1603602
- MAC move might be seen between the ICL and MC-LAG interface if adding or removing VLANs on the ICL interface. PR1605234
- On EX Series switches, the fxpc process might crash and generate a core file. PR1607372
- On EX4300 platform, the dcpfe process might crash and generate a core file. PR1608306
- DHCP packets might be received and then returned back to DHCP relay through the same interface on EX2300, EX3400, and EX4300 platforms. PR1610253
- Change in commit error message while configuring the same vlan-id with different vlan-name through openconfig CLI. PR1612566
- The non-VxLAN interfaces might not be able to egress packets on EX4400 line of switches involving with VxLAN configuration on VxLAN interfaces and VxLAN Network Identifier (VNI) configuration on non-VxLAN interfaces. PR1616683
- Due to the hardware programming error, CFM sessions fail. PR1619231
- On the EX4400 Series platform, EVPN type 5 routes might not be installed, when underlay is ECMP. PR1620808

Infrastructure

- On EX4600 platforms, the fxpc process might crash and generate core when router-advertisement-guard is configured under DHCP forwarding-options. PR1611480

Interfaces and Chassis

- On Junos OS platforms with VRRP failover-delay configured, changing VRRP mastership might cause peer device to re-learn VIP ARP entry on old primary interface because of the timing issue. PR1578126

J-Web

- J-Web allows a locally authenticated attacker to escalate their privileges to root. PR1594516

Junos Fusion Enterprise

- On all Junos OS and Junos Evo platforms with dual Routing Engine, when RE0 is reverted as primary Routing Engine, on rare occasions l2ald daemon might crash and cause an outage. PR1601817

Layer 2 Ethernet Services

- The DHCP client might be offline for about 120 seconds after sending the DHCPINFORM message. PR1587982

Platform and Infrastructure

- Broadcast traffic might be discarded when a firewall filter is applied to the loopback interface. PR1597548
- VLAN tagged traffic might be dropped with service provider style configuration. PR1598251
- On EX4300 platforms, when you configure `mac-move-limit` statement, forwarding the VRRP packets is not possible. PR1601005
- On EX4300 Series platforms, adding aggregated Ethernet configuration without child member might cause MAC/ARP learning issues. PR1602399
- ZTP does not work when you downgrade Junos OS Release 21.1R2.2 image to Junos OS Release 21.1R2.1 image. PR1603227

Routing Protocols

- The rpd core file might be observed because of the memory corruption. PR1599751
- The rpd process might crash and restart when NSR is enabled. PR1620463

Resolved Issues: 21.1R2

IN THIS SECTION

Class of Service (CoS)

- The buffer allocation for VCP ports might not get released in Packet Forwarding Engine after physically moving the port location. PR1581187

EVPN

- Traffic loss might be seen if aggregated Ethernet bundle interface with ESI is disabled on master Routing Engine followed by a Routing Engine switchover. PR1597300

General Routing

- MPPE-Send or Recv-key attribute is not extracted correctly by dot1xd. PR1522469
- FPC(s) might not boot-up on EX9214 in a certain condition. PR1545838
- Classifier is not programmed in the hardware and error logs might be seen in syslog. PR1548159
- "Cattle-Prod Daemon received unknown trigger (type Semaphore, id 1)" error messages seen on the vty when we issue CLI commands to fetch host route scale. PR1554140
- FPC with power related faults might get on-lined again after fabric healing has off-lined the FPC. PR1556558
- On the EX4300 device, script fails while committing the IPSec authentication configuration as the algorithm statement is missing. PR1557216

- The MAC addresses learned in a Virtual Chassis might fail due to aging out in the MAC scaling environment. PR1558128
- Some transmitting packets might get dropped due to the "disable-pfe" action is not invoked when the fabric self-ping failure is detected. PR1558899
- The tunable optics SFP+-10G-T-DWDM-ZR does not work. PR1561181
- EX3400VC - SMARTD pollutes syslog every 5 seconds after upgrading and rebooting the system. PR1562396
- On EX3400VC line of switches, the DAEMON-7-PVIDB throws syslog messages for every 12 to 14 minutes after you upgrade to Junos OS Release 19.1R3-S3. PR1563192
- The JWeb upgrade might fail on EX2300 and EX3400. PR1563906
- The DHCP client might not obtain IP address when dhcp-security is configured. PR1564941
- The new master Routing Engine post switchover might go into DB mode (or crash) on EX Series platforms. PR1565213
- On the EX Series Virtual Chassis, the following continuous message is observed: agentd-pfe-proxy_telemetry_publisher. PR1566528
- On the EX4600 line of switches, the following internal comment is displayed: Placeholder for QFX platform configuration. PR1567037
- The rpd crashes and generates a core file while booting the device. PR1567043
- The 40G DAC connection between EX9253 and the peers might not come up. PR1569230
- Packet loss might be observed when sample based action is used in firewall filter. PR1571399
- Port-mirroring might not work when the analyzer output is a trunk interface. PR1575129
- Protocol convergence between end nodes might fail when L2PT is enabled on transit switch. PR1576715
- The device implemented with different service image version might become VC member as unexpected. PR1576774
- MVR configuration cannot be configured on EX2300-C switches. PR1577905
- The fxpc process might crash on EX Series platforms. PR1578421
- Random or silent reboot might be seen on EX2300-24MP/EX2300-48MP platforms. PR1579576
- Some 40G ports might not be channelized successfully on the EX Series platforms. PR1582105

- The voice VLAN might not get assigned to the access interface. PR1582115
- When EX2300-MP in standalone mode is used as a DHCP server, initial set of packets received in the server might get dropped. PR1583983
- After performing NSSU, "timeout waiting for response from fpc0" error message is seen while checking version details. PR1584457
- DSCP rewriting might fail to work on EX2300. PR1586341
- The SNMP trap for MAC notifications might not be generated when an interface is added explicitly under switch-options. PR1587610
- Process dot1xd crash might be seen and re-authentication might be needed on EX9208 platform. PR1587837
- The rpd crash might be observed on the router running a scaled setup. PR1588439
- Packet loss might be observed on dynamically assigning VoIP VLAN. PR1589678
- Traffic loss might be observed for interface configured in subnet. PR1590040
- The LLDP packet might loss on the EX-4300MP platform if configuring LLDP on the management interface. PR1591387
- `show pfe filter hw` might generate an error "ERROR (dfw): Unknown group id: 21" message. PR1592096
- xSTP might not get configured when enabled on a interface with SP style configuration on all platforms. PR1592264
- Storm control profile might not be applied on EX2300 platforms. PR1594353
- On a EX4400 Virtual Chassis, log messages related to FAN settings will be observed in chassis traceoptions file. PR1594446
- On EX4400 Virtual Chassis, linecard member console might fail to redirect to Virtual Chassis master. PR1599625
- The upgrade using phone-home feature from Junos OS Release 20.2 or earlier to Junos OS Release 20.3 or later release will fail on EX2300/EX4650. PR1601722
- On EX4400 dot1x authentication might not work on EVPN/VXLAN enabled endpoints. PR1603015

Infrastructure

- EX 4300 VC/VCF : Observing HEAP malloc(0) detected. PR1546036

- For EX4400 product family, net installation (PXE) is not working. PR1577562
- Some MAC addresses might not be aged out on EX4300 platforms. PR1579293

Interfaces and Chassis

- MC-LAG interfaces might go down if the same VRRP group-id is configured on multiple IRB units.PR1575779
- The aggregated Ethernet interface might flap. PR1576533
- ARP resolution failure might occur during VRRP failover. PR1578126
- Incorrect advertisement threshold values are seen on VRRP groups when VRRP is configured on EX2300 switches. PR1584499

Layer 2 Features

- MAC addresses learnt from the MC-LAG client device might keep flapping between the ICL interface and MC-AE interface after one child link in the MC-AE interface is disabled. PR1582473

Layer 2 Ethernet Services

- Aggregated Ethernet interface flap might be seen during NSSU. PR1551925
- The DHCP client will be offline for 120 seconds after sending the DHCPINFORM message in the DHCP relay scenario. PR1575740

MPLS

- Incorrect EXP bit change might be seen in certain conditions under MPLS scenario. PR1555797

Platform and Infrastructure

- EX3400 VC - Console access on backup VC member is not allowed. PR1530106
- Upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart (vmcore). PR1557881
- The LLDP neighbor advertisement on EX4300 might send the incorrect 802.3 power format with TLV length 7 instead of length 12. PR1563105
- "Last flapped" timestamp for interface fxp0 gets reset every time "monitor traffic interface fxp0" is executed. PR1564323

- PFEX might crash when soft error recovery feature is enabled on a Packet Forwarding Engine. PR1567515
- On all EX9200 platforms with EVPN-VXLAN configured, the next hop memory leak in MX Series ASIC happens whenever there is a route churn for remote MAC-IP entries learned bound to the IRB interface in EVPN-VXLAN routing instance. When the ASIC's next hop memory partition is exhausted, the FPC might reboot. PR1571439
- Introduce two new major CMERRORs for XM chip-based line card to stabilize the running device. PR1574631
- DHCP packets with source IP as link-local address are dropped in EX4300. PR1576022
- The pfex might crash during PIC 4x 1G/10G SFP/SFP+ offline or online. PR1582457
- Firewall filter is not programmed correctly and traffic might be dropped unexpectedly. PR1586433
- The egress RACL firewall filter might not get programmed correctly on EX4300 platforms. PR1595797
- Broadcast traffic might be discarded when a firewall filter is applied to the loopback interface. PR1597548
- VLAN tagged traffic might be dropped with service provider style configuration. PR1598251
- The VRRP packets might not be forwarded when `mac-move-limit` statement is configured. PR1601005
- ZTP does not work when downgrading from Junos OS Release 21.1R2.2 image to Junos OS Release 21.1R2.1. PR1603227

Routing Protocols

- The untagged packets might not work on EX Series platforms. PR1568533
- BGP session carrying VPNv4 prefix with IPv6 next-hop might be dropped. PR1580578
- The rpd might crash in scaled routing instances scenario. PR1590638

Virtual Chassis

- EX4600/EX4300 mixed VC : Error message 'ex_bcm_pic_eth_uint8_set' is seen when changing configuration related to interface. PR1573173
- EX4300 VCP might not come up after upgrade when QSFP+-40G-SR4/QSFP+-40G-LR4/QSFP +40GE-LX4 is used. PR1579430

Resolved Issues: 21.1R1

Forwarding and Sampling

- Configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. PR1563641

General Routing

- While verifying the Last-change op-state value through XML, the rpc-reply message is inappropriate. PR1492449
- The mge interface might still stay up while the far end of the link goes down. PR1502467
- SNMP POE MIB walk produces either no results or sometimes results from the master Virtual Chassis whenever one of the Virtual Chassis is renamed. PR1503985
- The DHCP traffic might not be forwarded correctly when DHCP sends unicast packets. PR1512175
- Traffic loss might be observed on interfaces in a VXLAN environment. PR1524955
- On the EX2300, the following PoE message is observed: poe_get_dev_class: Failed to get PD class info. PR1536408
- EX4300-48MP : sFlow: dcpfe core file is generated when you use the `request chassis fpc slot slot_num restart` command. PR1536997
- On EX4300 platforms, the LLDP neighborship with the Voice over Internet Protocol (VoIP) phones cannot be established when LLDP is configured on the Power over Ethernet (PoE) enabled port on EX4300 and connects to the VoIP phone. PR1538482
- On the EX3400 and EX2300 switches, the upgrade fails due to the lack of available storage. PR1539293
- DHCP discover packet might be dropped if DHCP inform packet is received first. PR1542400
- The JNH memory might leak on the MX Series-based line cards PR1542882
- The Slaac-Snoopd child process generates a core file upon multiple switchovers on the Routing Engine. PR1543181
- In every software upgrade, host needs to get upgraded. PR1543890
- [Supportability] Improve Junos CLI outputs to display the host OS and kernel version in an easier and human-readable way. This enhancement is needed for all VMhost and Linux platforms. PR1543901
- The chip FPC might crash during the system booting. PR1545455

- Junos OS: Receipt of specific DHCPv6 packet may cause jdhcpd to crash and restart. PR1546166
- `show system software rollback` is not supported on EX4400-48F. PR1546605
- `show pfe route summary hw` shows random high free and 'Used' column for 'IPv6 LPM(< 64)' routes. PR1552623
- The configuration statement `action-shutdown` of storm control does not work for ARP broadcast packets. PR1552815
- On the EX9200 device, SF3 Fabric OIR issues is observed with Junos OS Release 23.1R1.8. PR1555727
- Traffic might be dropped when a firewall filter rule uses the then VLAN action. PR1556198
- On the EX4300 device, script fails while committing the IPsec authentication configuration as the algorithm statement is missing. PR1557216
- Observing error Opening configuration database: Could not open configuration database during usb upgrading. PR1561741
- The client authentication fails after GRES. PR1563431
- EX2300 shows high FPC CPU usage. PR1567438
- QinQ should not be licensed on EX2300, EX3400, EX4300, EX4300MP, and EX4400. PR1573179

Infrastructure

- On the EX4600 and EX4300 Virtual Chassis or Virtual Chassis Fabric, the VSTP configurations device goes unreachable and becomes nonresponsive after commit. PR1520351
- Error message during USB install: g_vfs_done():da0p1[READ(offset=65536, length=8192)]error = 5. PR1544736
- EX4300 Virtual Chassis or Virtual Chassis Fabric: Observing HEAP malloc(0) detected. PR1546036
- Traffic related to IRB interface might be dropped when mac-persistence-timer expires. PR1557229

Platform and Infrastructure

- DHCP binding does not happen after GRES. PR1515234
- lldp-receive-packet-count is not getting exchanged properly in l2pt operation for LLDP after configuring protocols. PR1532721

- On the EX4300 device, the LLDP neighborship might not come up with the non-aggregated Ethernet interfaces. PR1538401
- Junos OS: EX4300: FPC crash upon receipt of specific frames on an interface without L2PT or dot1x configured. PR1545530
- The targeted-broadcast feature might not work after a reboot. PR1548858
- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. PR1552603
- The targeted-broadcast feature might send out duplicate packets. PR1553070

Routing Protocols

- The OSPF neighborship gets stuck in the Start state after configuring the EVPN-VXLAN. PR1519244
- The OSPFv3 adjacency is not be established when IPsec authentication is enabled. PR1525870
- DCPFE crash might be observed while updating VRF for multicast routes during irb uninit. PR1546745
- The untagged packets might not work on QFX5000 platforms. PR1568533

User Interface and Configuration

- Removing the Flash component from the Monitor > Interfaces and DHCP pages removes other flash pages. PR1553176
- J-Web application package cannot be automatically updated for all the supported EX Series devices PR1563588

Documentation Updates

There are no errata and changes in Junos OS Release 21.1R3 documentation for the EX Series line of Switches.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for EX Series switches. 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 and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 7: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Junos OS Release Notes for JRR Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for the JRR 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for JRR Series Route Reflectors.

What's New in 21.1R3

There are no new features or enhancements to existing features in this release for JRR Series Route Reflectors.

What's New in 21.1R2

There are no new features or enhancements to existing features in this release for JRR Series Route Reflectors.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for JRR Series Route Reflectors.

Routing Protocols

- **Support for SR-IOV virtualization (JRR200)**—Starting in Junos OS Release 21.1R1, the JRR200 route reflector uses single-root I/O virtualization (SR-IOV) instead of full virtualization (emulated I/O) for network ports. SR-IOV helps to maximize the I/O throughput on the 10GbE SFP+ ports.

  SR-IOV improves:

- BGP convergence in a scaled environment.
- Throughput performance for BGP RIB sharding.

[See BGP sharding overview, rib-sharding, and update-threading.]

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for JRR Series Route Reflectors.

What's Changed in Release 21.1R3

There are no changes in behavior and syntax in Junos OS Release 21.1R3 for JRR Series Route Reflectors.

What's Changed in Release 21.1R2

There are no changes in behavior and syntax in Junos OS Release 21.1R2 for JRR Series Route Reflectors.

What's Changed in Release 21.1R1

There are no changes in behavior and syntax in Junos OS Release 21.1R1 for JRR Series Route Reflectors.

Known Limitations

There are no known limitations in hardware and software in this release for JRR Series Route Reflectors.

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

Open Issues

There are no open issues in hardware and software in this release for JRR Series Route Reflectors.

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

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for JRR 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: 21.1R3

IN THIS SECTION

General Routing

- JRR200, incorrect PEM load percentage for CLI `show chassis power`. PR1598728

Resolved Issues: 21.1R2

IN THIS SECTION

General Routing

- On JRR200, option-60 (Vendor-Class-Identifier) is not sent during ZTP. PR1582038

Resolved Issues: 21.1R1

IN THIS SECTION

General Routing

- The `request system power-off` and `request system halt` commands do not work as expected on JRR200. PR1534795
- Optics information of physical interfaces is not available for JRR200 on Junos OS. PR1537261

Documentation Updates

There are no errata or changes in Junos OS Release 21.1R3 documentation for JRR Series Route Reflectors.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for the JRR Series Route Reflector. 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 JRR200 Route Reflector Quick Start and Installation and Upgrade Guide.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 8: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Junos OS Release Notes for Juniper Secure Connect

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for Juniper Secure Connect. 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.

What's New

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for Juniper Secure Connect.

What's New in 21.1R3

There are no new features for Juniper Secure Connect in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features for Juniper Secure Connect in Junos OS Release 21.1R2.

What's New in 21.1R1

There are no new features for Juniper Secure Connect in Junos OS Release 21.1R1.

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for Juniper Secure Connect.

What's Changed in Release 21.1R3

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 21.1R3.

What's Changed in Release 21.1R2

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 21.1R2.

What's Changed in Release 21.1R1

There are no changes in behavior or syntax for Juniper Secure Connect in Junos OS Release 21.1R1.

Known Limitations

There are no known limitations for Juniper Secure Connect in Junos OS Release 21.1R3.

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

Open Issues

There are no open issues for Juniper Secure Connect in Junos OS Release 21.1R3.

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

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for Juniper Secure Connect.

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

Resolved Issues: 21.1R3

There are no resolved issues for Juniper Secure Connect in Junos OS Release 21.1R3.

Resolved Issues: 21.1R2

There are no resolved issues for Juniper Secure Connect in Junos OS Release 21.1R2.

Resolved Issues: 21.1R1

There are no resolved issues for Juniper Secure Connect in Junos OS Release 21.1R1.

Junos OS Release Notes for Junos Fusion for Enterprise

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for the Junos Fusion for enterprise. 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.

What's New

There are no new features or enhancements to existing features in Junos OS Releases 21.1R1, 21.1R2, or 21.1R3 for Junos fusion for enterprise.

What's Changed

There are no changes in behavior and syntax in Junos OS Releases 21.1R1, 21.1R2, or 21.1R3 for Junos fusion for enterprise.

Known Limitations

There are no known limitations in hardware or software in Junos OS Release 21.1R3 for Junos fusion for enterprise.

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

Open Issues

There are no known issues in hardware and software in Junos OS Release for 21.1R3 Junos fusion for enterprise.

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

Resolved Issues

There are no resolved issues in Junos OS Releases 21.1R1, 21.1R2, or 21.1R3 for Junos fusion for enterprise.

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

Documentation Updates

There are no corrections or changes in Junos OS Release 21.1R3 documentation for Junos fusion for enterprise.

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 for enterprise. Upgrading or downgrading Junos OS and satellite software might take several hours, depending on the size and configuration of the Junos fusion for 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/package-name.n.tgz
```

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

```
user@host> request system software add validate reboot source/package-name.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 for enterprise. See Configuring or Expanding a Junos fusion for enterprise.

For satellite device hardware and software requirements, see Understanding Junos fusion for 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 for 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 and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

  Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 9: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Downgrading Junos OS

Junos fusion for 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 for enterprise, follow the procedure for upgrading, but replace the `junos-install` package with one that corresponds to the appropriate release.

Junos OS Release Notes for Junos Fusion for Provider Edge

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for Junos Fusion for 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for Junos Fusion for Enterprise.

What's New in 21.1R3

There are no new features or enhancements to existing features in Junos OS Release 21.1R3 for Junos fusion for provider edge.

What's New in 21.1R2

There are no new features or enhancements to existing features for Junos fusion for provider edge in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for Junos Fusion for Provider Edge.

EVPN

- **Support for EVPN-MPLS (Junos fusion for provider edge)**—Starting in Junos OS Release 21.1R1, Junos fusion for provider edge supports EVPN-MPLS. EVPN-MPLS is a solution that extends Layer 2

VPN services over an MPLS network. Junos fusion for provider edge supports the connection of a customer edge (CE) device on the extended port of the satellite device in an EVPN-MPLS network.

[See Junos Fusion Provider Edge Supported Protocols.]

VPNs

- **Support for BGP MVPN (Junos fusion for provider edge)**—Starting in Junos OS Release 21.1R1, Junos fusion for provider edge supports BGP multicast VPN (MVPN). BGP MVPN is a method for implementing multiprotocol multicast services over a BGP MPLS Layer 3 VPN. Junos fusion for provider edge supports the connection of a BGP-based MVPN customer edge (CE) device on the extended ports of the satellite device in Junos fusion for provider edge.

  [See Junos Fusion Provider Edge Supported Protocols.]

- **Support for interprovider and carrier-of-carrier VPNs (Junos fusion for provider edge)**—Starting in Junos OS Release 21.1R1, Junos fusion for provider edge supports Interprovider and Carrier-of-Carrier VPNs. The Carrier-of-Carrier VPN service describes a hierarchical VPN (also known as a recursive VPN) model where one carrier (VPN service customer) transports its VPN traffic inside another carrier's VPN (VPN service provider). Junos fusion for provider edge currently supports provider edge (PE) routers for VPN service customers. In Junos OS Release 21.1R1, we introduce support for PE routers for VPN service providers along with VPN service customers.

  Interprovider VPNs provide connectivity between different service providers that are using separate autonomous systems (ASs) or one service provider that is using different ASs for different geographic locations. For Interprovider VPNs, Junos fusion for provider edge supports only intra-AS connection on an AS boundary router (ASBR) to the extended port.

  [See Junos Fusion Provider Edge Supported Protocols.]

What's Changed

There are no changes in behavior and syntax in Junos OS Releases 21.1R1, 21.1R2, or 21.1R3 for Junos fusion for provider edge.

Known Limitations

There are no known limitations in hardware or software in Junos OS Release 21.1R3 for Junos fusion for provider edge.

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

Open Issues

There are no known issues in hardware or software in Junos OS Release 21.1R3 for Junos fusion for provider edge.

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

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for Junos fusion for provider edge.

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

Resolved Issues: 21.1R3

IN THIS SECTION

Learn about the issues fixed in these releases for Junos fusion for provider edge.

Junos Fusion Provider Edge

- Configuring the port mirroring firewall filter in a bridge domain with IRB might cause traffic loss over IRB. PR1607750

Resolved Issues: 21.1R2

IN THIS SECTION

Interfaces and Chassis

- On MX platforms in Junos fusion scenario, if targeted-distribution is configured for aggregated Ethernet, vlan-demux, or PPPoE interfaces whose underlying legs are on FPC numbers greater than 32 (for example, ge-101/0/0), then the dcd process might crash and FPC might be stuck in ready state. PR1601566
- On Junos Fusion system with MX as aggregation devices, the 100 G aggregated Ethernet interfaces might flap upon unrelated configuration changes. PR1602656

Resolved Issues: 21.1R1

There are no fixed issues in the Junos OS Release 21.1R1 for Junos fusion for provider edge.

Documentation Updates

There are no corrections or changes in Junos OS Release 21.1R3 documentation for Junos fusion for provider edge.

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 for 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 Installation and Upgrade Guide.

The download and installation process for Junos OS Release 21.1R3 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 use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot source/jinstall64-21.1R3.SPIN-domestic-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot source/jinstall-21.1R3.SPIN-domestic-signed.tgz
```

All other customers, use the following commands.

- For 64-bit software:

> **NOTE**: We recommend that you use 64-bit Junos OS software when implementing Junos fusion for provider edge.

```
user@host> request system software add validate reboot source/jinstall64-21.1R3.SPIN-export-signed.tgz
```

- For 32-bit software:

```
user@host> request system software add validate reboot source/jinstall-21.1R3.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 21.1R3 `jinstall` package, you cannot return to the previously installed software by issuing the `request system software rollback` command. Instead, you must issue the `request system software add validate` command and specify the `jinstall` package that corresponds to the previously installed software.

Upgrading an Aggregation Device with Redundant Routing Engines

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

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

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

Preparing the Switch for Satellite Device Conversion

Satellite devices in a Junos fusion topology use a satellite software package that is different from the standard Junos OS software package. Before you can install the satellite software package on a satellite device, you first need to upgrade the target satellite device to an interim Junos OS software version that

can be converted to satellite software. For satellite device hardware and software requirements, see Understanding Junos fusion Software and Hardware Requirements

NOTE: The following conditions must be met before a standalone switch that is running Junos OS Release 14.1X53-D43 can be converted to a satellite device when the action is initiated from the aggregation device:

- The switch can be converted to only SNOS 3.1 and later.
- Either the switch must be set to factory-default configuration by using the `request system zeroize` command, or the following command must be included in the configuration: `set chassis auto-satellite-conversion`.

Customers with EX4300 switches, use the following command:

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

Customers with QFX5100 switches, use the following command:

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

When the interim installation has completed and the switch is running a version of Junos and OS on one line 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 for 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 21.1R3, you must also upgrade your satellite device to Satellite Device Software version 3.1R1.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

  Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 10: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Downgrading from Junos OS Release 21.1

To downgrade from Release 21.1 to another supported release, follow the procedure for upgrading, but replace the 21.1 `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.

Junos OS Release Notes for MX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for the MX Series routers.

What's New in 21.1R3

There are no new features or enhancements to existing features in Junos OS Release 21.1R3 for MX Series routers.

What's New in 21.1R2

There are no new features or enhancements to existing features for MX Series routers in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for the MX Series routers.

Hardware

- We've added the following features to the MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) and MPC11E (MX2K-MPC11E) in Junos OS Release 21.1.

Table 11: Feature Support on MPC10E and MPC11E on MX Series Routers

Feature	Description
EVPN	• Configure inner source MAC address for flexible VXLAN tunnels—Use the Juniper Extension Toolkit (JET) RIB Service API to configure the source MAC address used in IPv4 and IPv6 flexible VXLAN tunnel encapsulation profiles. If you don't specify a source MAC address, the default source MAC address 00:00:5e:00:52:01 is used to encapsulate IPv4 and IPv6 flexible VXLAN tunnels. [See Understanding Programmable Flexible VXLAN Tunnels and Juniper Extension Toolkit (JET).] • Support for auto-derived route targets on EVPN-MPLS.Junos OS supports the automatic derivation of route targets on EVPN-MPLS in an MPC10E line card on an MX Series router. When you enable the auto-derived route target feature, route targets are automatically derived from the VLAN ID for EVPN Type 2 and EVPN Type 3 routes and can be imported to the EVPN routing instance table. To enable the auto-derived route targets option, include the `auto` statement at the [`edit routing-instances routing-instance-name protocols evpn vrf-target`] hierarchy level. [See Auto-derived Route targets.] • Support for IPv4 unicast VXLAN encapsulation optimization on MPC10E and MPC11E line cards running on MX240, MX480, MX960, MX2008, MX2010, and MX2020 routers. By default, these routers optimize VXLAN-encapsulated throughput for IPv4 unicast packets that are 512 through 1500 bytes in size over the following VXLAN tunnel types: • PIM-based VXLAN • EVPN-VXLAN • Static VXLAN This feature doesn't provide additional optimization over EVPN Type 5 tunnels (which are already optimized), and is not supported with forwarding table filters. [See Understanding VXLANs.]

Table 11: Feature Support on MPC10E and MPC11E on MX Series Routers *(Continued)*

Feature	Description
High availability (HA) andresiliency	• MX Series Virtual Chassis (MX-VC) support for MPC10E-10C-MRATE and MPC10E-15C-MRATE (MX240, MX480, and MX960)—You can operate the MPC10E-10C-MRATE and MPC10E-15C-MRATE line cards in a router in an MX Series Virtual Chassis. The MPC10E support in MX-VC is only for uplink usage. [See Virtual Chassis Components Overview.]
Juniper Extension Toolkit (JET)	• Support for static backup paths with IP-in-IP tunnel encapsulation and provisioning APIs (MX240, MX480, MX960, MX2010 and MX2020)—We've enhanced Juniper Extension Toolkit (JET) APIs to enable a controller to set up underlay network backup paths that use IP-in-IP tunnels with IPv4 encapsulation. [See Juniper Extension Toolkit (JET).]
Layer 2 features	• Support for MAC statistics (MX-Series)— You can enable MAC statistics for Layer 2 traffic on MPC10E-15C-MRATE, MPC10E-10C-MRATE, and MX2K-MPC11E MPC line cards. To enable MAC statistics at the bridge domain, include the `mac-statistics` configuration statement at the [`edit bridge-domains <bridge-domain name> bridge-options`] hierarchy level. To enable MAC statistics at the global level, you need to include the `global-mac-statistics` configuration statement at the [`edit protocols l2-learning`] hierarchy level. [See mac-statistics and global-mac-statistics.] • Support for Multiple VLAN Registration Protocol (MVRP) and Ethernet ring protection switching (ERPS). [See Understanding Multiple VLAN Registration Protocol (MVRP) for Dynamic VLAN Registration and Ethernet Ring Protection Switching Overview.]

Table 11: Feature Support on MPC10E and MPC11E on MX Series Routers *(Continued)*

Feature	Description
Port security	• Support for Media Access Control Security (MACsec) on logical interfaces (MPC10E and MPC11E). VLAN tags are transmitted in clear text, which allows intermediate switches that are MACsec-unaware to switch the packets based on the VLAN tags. [See Media Access Control Security (MACsec) over WAN.]

Table 11: Feature Support on MPC10E and MPC11E on MX Series Routers *(Continued)*

Feature	Description

Services applications	• Support for Mapping of Address and Port with Encapsulation (MAP-E) and inline 6rd (MPC10E and MX2K-MPC11E)— You can configure MAP-E and inline IPv6 rapid deployment (inline 6rd) on the following MPCs:
	  • MPC10E-15C-MRATE and MPC10E-10C-MRATE on MX240, MX480, and MX960 routers
	  • MX2K-MPC11E on MX2010 and MX2020 routers
	  [See Configuring Mapping of Address and Port with Encapsulation (MAP-E) and Configuring Inline 6rd.]
	• Support for tunnel interfaces on the MPC10E line card—Junos OS supports three tunnel interfaces on the MPC10E line card: generic routing encapsulation (GRE) tunnel, logical tunnel (LT), and virtual tunnel (VT).
	  • The GRE tunnel interface supports the `tunnel` statement with these options: `destination`, `key`, `source`, `traffic-class` and `ttl`. The `copy-tos-to-outer-ip-header` statement is also supported.
	  • The LT interface supports the family `inet`, `inet6`, and `iso` options. The `encapsulation` statement supports the Ethernet and VLAN physical interface options only.
	  • The VT interface supports the family `inet` option only.
	  [See Tunnel Services Overview].
	• AMS support (MX240, MX480, MX960, MX2010, and MX2020 routers)—Junos OS supports aggregated multiservices (AMS) interfaces on the MPC10E and MX2K-MPC11E line cards to provide load balancing and high availability features for stateful firewall and NAT services. You can configure AMS interfaces with next-hop style service sets and with MS-MPC or MS-MIC only.
	  [See Understanding Aggregated Multiservices Interfaces.]

Table 11: Feature Support on MPC10E and MPC11E on MX Series Routers *(Continued)*

Feature	Description
System management	• Support for Synchronous Ethernet over link aggregation group interfaces (MX240, MX480, and MX960)—MPC10E line cards support Synchronous Ethernet over a link aggregation group (LAG). [See Synchronous Ethernet Overview.] • Support for PTP over Ethernet, hybrid mode, and G.8275.1 profile (MX240, MX480, and MX960)—MPC10E line cards support Precision Time Protocol (PTP) over Ethernet, G.8275.1 profile, and hybrid mode. [See Precison Time Protocol Overview and Understanding Hybrid Mode.] • Support for PTP over Ethernet and hybrid mode over link aggregation group interfaces (MX240, MX480, and MX960)—MPC10E line cards support Precision Time Protocol (PTP) over Ethernet and hybrid mode over a link aggregation group (LAG). [See Understanding Hybrid Mode and Precison Time Protocol Overview.]

Dynamic Host Configuration Protocol

- **Include DHCP option 61 in Radius Access Request (MX240, MX480, and MX960 routers)**—Starting in Junos OS Release 21.1R1, you can configure DHCP to use the client identifier (DHCP Option 61) in the username that is passed to the external AAA authentication service when the DHCP client logs in. You can also configure options to exclude headers and to use automatic ASCII hex encoding to obtain the preferred string for authentication. This feature is supported for DHCP server and relay in DHCPv4, DHCPv6, and dual stack.

  [See Creating Unique Usernames for DHCP Clients.]

EVPN

- **Flow-aware transport pseudowire support for EVPN-VPWS (MX Series routers and EX9200 switches)**—Starting in Junos OS Release 21.1R1, you can statically configure provider edge (PE) devices to use flow-aware transport (FAT) pseudowire labels in an EVPN virtual private wire service (VPWS) routing instance with an IP/MPLS underlay fabric. PE devices use these labels to load-

balance EVPN-MPLS packets across ECMP paths or link aggregation groups (LAGs) without needing to do deep packet inspection of the payload.

To enable FAT pseudowire load balancing in an `evpn-vpws` routing instance:

- Configure `flow-label-transmit-static` on PE devices to insert FAT flow labels into VPWS pseudowire packets sent to remote PE devices.
- Configure `flow-label-receive-static` on PE devices to remove FAT flow labels from VPWS pseudowire packets received from remote PE devices.

You can configure these statements for all pseudowires in the routing instance or for pseudowires associated with a specific interface in the routing instance.

[See FAT Flow Labels in EVPN-VPWS Routing Instances, flow-label-receive-static, and flow-label-transmit-static.]

Interfaces

- **Support for Q-DD-400G-LR4-10 and Q-DD-4X100G-LR optics on MPC11E line card (MX2010 and MX2020)**—Starting in Junos OS Release 21.1R1, you can use the Q-DD-400G-LR4-10 and Q-DD-4X100G-LR 400GE optics on the MPC11E line card in the MX2010 and MX2020 routers.

  [See Hardware Compatibility Tool.]

- **Support for DDOS telemetry (MX Series)**–Starting in Junos OS Release 21.1R1, you can get the DDOS telemetry statistics with the help of Jvision from MPC1, MPC2, MPC3, MPC5, MPC6, MPC7, MPC8, and MPC9 line cards. Use configuration statements:

  - `[show services analytics sensor ddos]` - to see the DDOS configuration
  - `[show agent sensors]` - to see the details about configured sensors

  [See sensor (Junos Telemetry Interface), Guidelines for gRPC and gNMI Sensors (Junos Telemetry Interface), show agent sensors, and show configuration services analytics sensor ddos.]

Junos Telemetry Interface

- **VCP interchassis link port statistics and optimized queue statistics for aggregated Ethernet bundle support with JTI (MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX100016, and vMX)**—Starting in Junos OS Release 21.1R1, you can use Junos telemetry interface (JTI) with remote procedure call (gRPC) services to export Virtual Chassis port (VCP) interchassis link statistics and queue statistics. A Virtual Chassis operator can subscribe to supported sensors to monitor the health of MX-VC interchassis links and be able to diagnose and solve issues that may arise when links become unhealthy. This feature also optimizes queue statistics collection, limiting the data and statistical polling to queues that are part of an

interface set for which there is an active telemetry subscription. Doing so reduces resources and creates a more stable operating environment.

The supported new sensors are:

- Virtual Chassis port sensor, which provides interchassis links basic error and operational state for the interface (resource path **/junos/system/mxvc/members/member[memberID=<ID>]/virtual-chassis-ports/virtual-chassis-port[vcp-interface-name=<vcp-interface-port-string>]**). This sensor collects the same data displayed when using the operational mode command `show interfaces vcp-x/x/x extensive` and `show virtual-chassis vc-port`.

- Virtual Chassis heartbeat sensor, which provides insight into the overall link health of the Virtual Chassis system (resource path **/junos/system/mxvc/members/member[memberID=<ID>]/heartbeat**). The Virtual Chassis heartbeat sensor, when active, periodically sends and receives information between the chassis to keep track of the connection state. This sensor collects the same data displayed when using the operational mode command **show virtual-chassis heartbeat detail**.

- Virtual Chassis high/low DDoS queue statistics sensor (resource path **/junos/system/mxvc/members/member[memberID=<ID>]/ddos-protocols/ddos-protocol[packetType=<packet_type>]/**). This sensor tracks high and low statistics used in distributed denial-of-service (DDoS) protection. This sensor is available on the global primary chassis. No sensor output is collected for a backup chassis. This sensor collects the same data displayed when using the operational mode command **show ddos-protection protocols virtual-chassis statistics terse**.

- Timestamp sensor (resource path **/junos/system/subscriber-management/dynamic-interfaces/interface-sets/queue-statistics/interface-set[container-index=<container-index>]/fpcs/fpc[slot=<slot>]/last-update-time** and **/junos/system/subscriber-management/dynamic-interfaces/interfaces/queue-statistics/interface[sid=<session-identifier>]/fpcs/fpc[slot=<slot>]/last-update-time**). The timestamp sensor keeps track of when the FPC calculated the queue statistics for a given FPC leg. This provides a better way to track the validity of the data and is a basis for more accurate rate calculations.

The enhanced sensors are:

- Per-subscriber queue statistics sensor (resource path **/junos/system/subscriber-management/dynamic-interfaces/interfaces/queue-statistics/interface[sid=<session-identifier>]/fpcs/fpc[slot=<slot>]/queues/queue[queue-no=<queue-index>]** and **/junos/system/subscriber-management/dynamic-interfaces/interface-sets/queue-statistics/interface-set[container-index=<container-index>]/fpcs/fpc[slot=<slot>]/queues/queue[queue-no=<queue-index>]**). This sensor can include the additional leafs `current-polling-interfaces` and `current-polling-interface-sets`. When you provide a corresponding index (SID for interface or container ID for interface set), polling for queue statistics is enabled only for that corresponding index. If no index is specified, polling all indexes queue statistics is enabled.

[See Junos Telemetry Interface User Guide.]

- **Optimized chassis sensor support with JTI (MX2010 and MX2020 with MPC11E and MPC9E line cards)**—Starting in Junos OS Release 21.1R1, you can use Junos telemetry interface (JTI) with remote procedure call (gRPC) services to export additional chassis sensors from an MX2010 or MX2020 router to an outside collector.

  For network debugging, there are system-generated logs and SNMP traps available. However, some parameters, such as power entry module (PEM) voltage and PEM supply failure, have not been available in telemetry. We've now introduced these additional system parameters through chassis sensors that support messages logged as part of the SNMP traps for a field-replaceable unit (FRU), fan, PEM, or plane.

  Use the resource path **/components/component/properties/property** in subscriptions for these additional chassis sensors:

  - power-supply-failed
  - chassisd-pem-breaker-trip
  - chassisd-pem-voltage
  - fan-blower-removed
  - pem-not-powered
  - chassisd-zone-blowers-speed-type
  - chassisd-zone-blowers-speed
  - temprature-back-to-normal
  - over-temprature
  - fru-failed
  - plane-fru-check
  - plane-online

  [See Junos Telemetry Interface User Guide.]

MPLS

- **Nonstop active routing (NSR) support for controller-initiated RSVP label-switched paths (LSPs) (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 21.1R1, we support NSR for controller-initiated RSVP-based point-to-point (P2P) and point-to-multipoint (P2MP) LSPs. The primary Routing Engine synchronizes all RSVP LSPs initiated by Path Computation Elements (PCEs),

including multicast flow specifications for any PCE-initiated P2MP LSPs, with the backup Routing Engine. This ensures zero traffic loss for the traffic carried over PCE-initiated RSVP LSPs during Routing Engine switchovers. This feature is enabled when NSR is configured.

[See PCEP Configuration.]

- **BGP Classful Transport planes (BGP-CT) to facilitate service mapping over colored tunnels (ACX Series, PTX Series, MX Series)**—Starting in Junos OS Release 21.1R1, you can classify colored transport tunnels (RSVP, IS-IS flexible algorithm) in your network into transport classes and map service routes over an intended transport class. You can also extend the transport tunnels to span across multiple domains (ASs or IGP areas) by using the new BGP transport address family called BGP Classful Transport (BGP CT).

  This feature lays the foundation for network slicing and allows the different domains to interoperate irrespective of the transport signaling protocols used in each domain.

  [See BGP Classful Transport Planes Overview.]

- **Install prefixes for RSVP-TE LSPs using PCEP (MX Series, PTX Series, QFX Series)**–Starting in Junos OS Release 21.1R1, you can configure different prefixes for Path Computation Element (PCE)-initiated and PCE-delegated RSVP-TE LSPs using the Path Computation Element Protocol (PCEP). Prior to this feature, for PCE-initiated LSPs, you could install prefixes as routes through templates and map the templates to the LSPs. For Path Computation Client (PCC)-configured LSPs, although you could install prefixes on the device, this information was not reported to the PCE.

  With this feature, you can install prefixes for external RSVP-TE LSPs through PCEP communication, and enable the PCC to report installed prefixes for all local RSVP-TE LSPs to the PCE . This support provides you better traffic engineering capabilities and allows Junos OS to interoperate with other vendor's PCC or PCE.

  [See PCEP Overview.]

Multicast

- **Controller-based BGP multicast signaling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 21.1R1, we've introduced controller-based BGP multicast signaling that can be used instead of hop-by-hop signaling to program multicast forwarding states on routers. An external controller that is aware of the topology and network events within that topology calculates the optimum multicast trees between the source and receivers. The external controller then uses BGP signaling to send a new type of BGP network layer reachability information (NLRI) with modified attributes to convey the multicast state information to all the routers on the multicast trees.

  You can use this feature instead of multicast routing protocols, such as Protocol Independent Multicast (PIM) or multipoint LDP (MLDP). You can enable this feature using `bgpmcast` configuration option at the `[edit protocols]` hierarchy.

- **MVPN live-live solution support (MX Series)**—Starting in Junos OS Release 21.1R1, we've added support to enable the MVPN live-live feature in next-generation multicast VPN (MVPN) with multicast LDP point-to-multipoint (P2MP) provider tunnel. This feature helps to keep your network live all the time.

  To enable the MVPN live-live solution:

  - Configure the `sender-based-rpf` option by running the `set routing-instances` *`routing-instance-name`* `protocols mvpn sender-based-rpf` command. This option is disabled by default.
  - Configure the `hot-root-standby` option by running the `set routing-instances` *`routing-instance-name`* `protocols mvpn hot-root-standby` command. You can configure this option only if sender-based RPF is enabled.

  When you enable this configuration, the receiving PE automatically switches over to the backup path if it encounters any failure while forwarding the traffic from the primary path to the customer network. The transition from primary path to backup path happens in less than 50 milliseconds.

  For previous Junos OS releases, we provided support only for RSVP-TE and IR provider tunnels.

  [See sender-based-rpf and hot-root-standby.]

Network Management and Monitoring

- **Ephemeral configuration database support for load update operations (EX9200, MX5, MX10, MX80, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMX)**—Starting in Junos OS Release 21.1R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database using a load update operation. To perform a load update operation, set the `<load-configuration> action` attribute to `update`.

  [See <load-configuration>.]

- **Ephemeral configuration database support for synchronous commit synchronize operations on dual Routing Engine devices (MX240, MX480, MX960, MX2010, MX2020, MX10003, MX10008, and MX10016)**—Starting in Junos OS Release 21.1R1, you can configure the ephemeral database to execute commit synchronize operations using a synchronous commit model on dual Routing Engine devices. The synchronous commit model enables you to reliably use the ephemeral database on devices that have graceful Routing Engine switchover (GRES) or non-stop routing (NSR) enabled. To use the synchronous commit model for the ephemeral database, configure the `commit-synchronize-model synchronous` statement at the `[edit system configuration-database ephemeral]` hierarchy level.

  [See Understanding the Ephemeral Configuration Database.]

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020,**

MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

[See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

- **SNMP support for carrier-grade NAT PBA monitoring (MX Series)**—Starting in Junos OS Release 21.1R1, you can get port block allocation (PBA) information about MS-MPC and unified services framework (USF)MX-SPC3 - related aspects using two new MIB objects and two new MIB tables:

  - New MIB object jnxNatSrcNumAddressMapped under the MIB table jnxSrcNatStatsTable, and a new MIB table jnxNatPbaStatsTable to get information about MS-MPC-PIC and MS-MIC

  and

  - New MIB object jnxJsNatSrcNumAddressMapped under the MIB table jnxJsSrcNatStatsTable,and a new MIB table jnxJsNatPbaStatsTable to get information about MX-SPC3.

  [See SNMP MIBs and Traps Supported by Junos.]

- **sFlow support for IP-IP traffic (MX240, MX480, and MX960)**—Starting in Junos OS Release 21.1R1, you can use sFlow technology to sample egress sFlow for IP over IP (IP-IP) traffic at the tunnel entry point, transit device, and tunnel endpoint on a physical port. sFlow sampling is supported for IP-IP tunnels with an IPv4 outer header that carry IPv4 or IPv6 traffic. Tunnel header encapsulation is done by either dynamic tunnel or FTI (Flexible Tunnel Interface). You can use sFlow monitoring technology to randomly sample network packets from IP-IP tunnels and to send the samples to a destination collector for monitoring.

  [See Overview of sFlow Technology and Configuring IP Tunnel Interfaces.]

- **HMAC-SHA-2 authentication protocol support for users of SNMPv3 USM (MX Series and SRX Series)**—Starting in Junos OS Release 21.1R1, you can configure HMAC-SHA-2 authentication protocols for users of the SNMPv3 user-based security model (USM) with the following new CLI configuration statements:

  - authentication-sha224
  - authentication-sha256
  - authentication-sha384

- authentication-sha512

We've introduced these statements for local-engine users at `[edit snmp v3 usm local-engine user username]` and for remote-engine users at `[set snmp v3 usm remote-engine engine-id user user-name]`.

[See authentication-sha224, authentication-sha256, authentication-sha348, and authentication-sha512.]

OpenConfig

- **OpenConfig support for VLAN interfaces (MX240)**—Junos OS Release 21.1R1 supports the following OpenConfig Data Model openconfig-interfaces.yang version 2.4.3 files for VLAN interfaces:
  - **/interfaces/interface/subinterfaces/subinterface/oc-vlan:vlan/oc-vlan:match/oc-vlan:single-tagged/oc-vlan:config/oc-vlan:vlan-id**
  - **/interfaces/interface/subinterfaces/subinterface/oc-vlan:vlan/oc-vlan:match/oc-vlan:double-tagged/oc-vlan:config/oc-vlan:inner-vlan-id**
  - **/interfaces/interface/subinterfaces/subinterface/oc-vlan:vlan/oc-vlan:match/oc-vlan:double-tagged/oc-vlan:config/oc-vlan:outer-vlan-id**

  [See Mapping OpenConfig VLAN Commands to Junos Configuration.]

- **OpenConfig support for GRE tunnel interfaces (MX5, MX10, MX40, MX150, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMX)**—Junos OS Release 21.1R1 supports the following OpenConfig Data Model openconfig-if-tunnel.yang version 0.1.1 sensors:
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:config/oc-tun:src**
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:config/oc-tun:dst**
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:config/oc-tun:ttl**
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:ipv4/oc-tun:addresses/oc-tun:address/oc-tun:config/oc-tun:ip**
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:ipv4/oc-tun:addresses/oc-tun:address/oc-tun:config/oc-tun:prefix-length**
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:ipv4/oc-tun:config/oc-tun:mtu**
  - **/interfaces/interface/oc-tun:tunnel/oc-tun:config/oc-tun:gre-key**

  This feature includes converting OpenConfig configuration schemas to Junos OS configuration schemas. For example, the OpenConfig command`set openconfig-interfaces:interfaces interface gr-0/0/0`

`openconfig-if-tunnel:tunnel config src 10.1.1.1` maps to the Junos OS command `set interfaces gr-0/0/0 unit 0 tunnel source 10.1.1.1`.

[See OpenConfig User Guide.]

Platform and Infrastructure

- **Next Gen Services (MX240, MX480, and MX960 with MX-SPC3)**— Starting in Junos OS Release 21.1R1, we support IPsec (a Next Gen Services component) on the listed MX Series routers with the MX-SPC3 services card installed. To configure IPsec on MX Series routers with MX-SPC3, use the CLI configuration statements at the [`edit security`] hierarchy level. On MX Series routers with MS-MPC/MS-MIC line cards, you configure the feature at the [`edit services`] hierarchy level.

  NOTE: MX240, MX480, and MX960 routers with MS-MPC/MS-MIC and MX-SPC3 support Next Gen Services. We introduced this support in Junos OS Release 19.3R2.

  Table 12: Next Gen Services Supported on MX-SPC3

Feature	Description
MX-SPC3 IPsec VPN Feature License	You require a valid license to use the IPsec VPN feature on your MX Series devices with the MX-SPC3 services card. This is a binary license. The `show system license` command output displays the license count as 0 when no license is installed and 1 when a valid license is installed. You won't be able to establish IPsec VPN tunnels if you don't have a valid license to use the feature. However, tunnels that are currently active will continue to stay up if your license expires. You cannot reestablish IPsec VPN tunnels that go down after the expiry of the license until you install a valid license. See Managing Licenses.
IPsec VPN	The MX-SPC3 services card provides consistent IPsec VPN capability across security and routing platforms. You configure IPsec for the MX-SPC3 at the [`edit security`] hierarchy level. See Next Gen Services Overview

Table 12: Next Gen Services Supported on MX-SPC3 *(Continued)*

Feature	Description
AutoVPN preshared key (PSK) on MX-SPC3	To allow different IKE preshared keys used by the VPN gateway to authenticate the remote peer, use our new CLI statements `seeded-pre-shared-key ascii-text` or `seeded-pre-shared-key hexadecimal` at the `[edit security ike gateway gateway_name]` hierarchy level. To allow the same IKE preshared key used by the VPN gateway to authenticate the remote peer, use the existing CLI command `pre-shared-key ascii-text` or `pre-shared-key hexadecimal`. During authentication of the remote peer, use the `general-ikeid` statement at the `[edit security ike gateway gateway_name dynamic]` hierarchy level to bypass the IKE-ID validation. See AutoVPN on Hub-and-Spoke Devices.
Add new members to existing aggregated multiservice (AMS) bundle for IPsec service	To add new members to an AMS bundle (for IPsec services) without impacting the traffic on the existing AMS bundle, configure the `no-bundle-flap` statement under the `[edit interfaces interface-name load-balancing-options]` hierarchy in non-HA mode. During the configuration change, the existing members in the AMS bundle don't flap. See Understanding Aggregated Multiservices Interfaces for Next Gen Services.
PowerMode IPsec	The MX-SPC3 card supports PowerMode IPsec (PMI) with vector packet processing (VPP) and Intel Advanced Encryption Standard New Instructions (AES-NI), leading to IPsec performance improvements. You can enable PMI processing by using the `set security flow power-mode-ipsec` command. To disable PMI processing, use the `delete security flow power-mode-ipsec` command. MX-SPC3 also supports the fat tunnel feature that improves the performance of a single tunnel. If one of the tunnels is loaded with traffic and other tunnels have less traffic, the resources are shared within the fat group. This results in an even CPU utilization of the resources. To enable this feature, configure the `fat-core` statement at the `[edit security distribution-profile]` hierarchy level. You must configure the PMI feature first to enable the fat tunnel feature. See Improving IPsec Performance with PowerMode IPsec, Understanding Symmetric Fat IPsec Tunnel, and power-mode-ipsec.

Table 12: Next Gen Services Supported on MX-SPC3 *(Continued)*

Feature	Description
Support for mobility in CGNAT–XLAT464	We've upgraded the current dual-translation (464XLAT) feature by introducing `clat-ipv6-prefix-length` at the source NAT rule hierarchy level. You can use a single NAT rule with this configuration parameter in place of multiple source NAT rules with different source-address and customer-side translator (CLAT)-prefix values. This simplifies the configuration method for certain use case scenarios.
Support for time zones in carrier-grade NAT	Support for syslog timestamp (local system time stamp) using the `utc-timestamp` statement at the [`edit interfaces interface-name services-options`] hierarchy level.
Network Address Translation - Port Translation (NAT-PT)	We support NAT-PT with the DNS ALG service on the MX-SPC3 services card. See Configuring the DNS ALG.
MPC10E interoperability	The MPC10E (MPC10E-15C-MRATE and MPC10E-10C-MRATE) line card interoperates with the MX-SPC3 services card to support the NAT and stateful firewall Layer 3 services. See Protocols and Applications Supported by MX-SPC3 Services Card

[See Next Gen Services Overview.]

Port Security

MACsec bounded delay protection (MX10003 routers)—Starting in Junos OS Release 21.1R1, you can configure bounded delay protection on MX10003 routers. MACsec bounded delay protection prevents the delivery of a frame when the frame is delayed by two seconds or longer. This feature enables the detection of delayed MACsec frames that result from a man-in-the-middle attack.

Routing Protocols

- **IS-IS link delay measurement and advertising (MX Series)**—Starting in Junos OS Release 21.1R1, you can measure and advertise various performance metrics in IP networks with scalability, by using several IS-IS probe messages. These metrics can then be used to make path-selection decisions based on network performance.

  [See How to Enable Link Delay Measurement and Advertising in IS-IS, delay-measurement, and delay-metric.]

- **Support for configuring multiple independent IGP instances of IS-IS (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 21.1R1, you can configure and run multiple independent IGP instances of IS-IS simultaneously on a router.

  NOTE: Junos OS does not support configuring the same logical interface in multiple IGP instances of IS-IS.

  [See How to Configure Multiple Independent IGP Instances of IS-IS.]

- **Support for BGP Auto-discovered Neighbor (MX Series, PTX1000, PTX10008, QFX5120-32C, QFX5200, QFX5210, and QFX10008)**—Starting in Junos OS Release 21.1R1, we support BGP auto-discovered neighbors using IPv6 Neighbor Discovery Protocol (ND). With this feature, you can enable BGP to create peer neighbor sessions using link-local IPv6 addresses of directly connected neighbor devices. You need not specify remote or local neighbor IP addresses.

  To enable peering for a given interface or set of interfaces without specifying the local or remote neighbor addresses, configure the `peer-auto-discovery` statement at the [`edit fabric protocols bgp group <name> dynamic-neighbor <name>`] hierarchy level.

  [See BGP Auto-Discovered Neighbors, and peer-auto-discovery.]

Segment Routing

- **Avoid microloops in IS-IS-SRv6 networks (MX Series with MPC7E, MPC8E and MPC9E line cards)** — Starting in Junos OS Release 21.1R1, you can enable post-convergence path calculation on a device to avoid microloops if a link or metric changes in an SRv6 network. Note that microloop avoidance is not a replacement for local repair mechanisms such as topology-independent loop-free alternate (TI-LFA), which detects local failure very fast and activates a precomputed loop-free alternative path.

  To configure microloop avoidance in an SRv6 network, include the `microloop avoidance post-convergence-path delay` *`milliseconds`* statement at the [`edit protocols isis spf-options`] hierarchy level.

  [See How to Configure Microloop Avoidance for IS-IS in SRv6 Networks.]

- **SRv6 network programming in IS-IS (MX Series with MPC10 and MPC11 line cards)**—Starting in Junos OS Release 21.1R1, you can configure segment routing in a core IPv6 network without an MPLS data plane. This feature is useful for service providers whose networks are predominantly IPv6 and have not deployed MPLS. Such networks depend only on the IPv6 headers and header extensions for transmitting data. This feature also benefits networks that need to deploy segment routing traffic through transit routers that do not have segment routing capability yet. In such networks, the SRv6 network programming feature can provide the flexibility to leverage segment routing without deploying MPLS.

To enable SRv6 network programming in an IPv6 domain, include the `srv6` statement at the`[edit routing-options source-packet-routing]` hierarchy level.

To advertise the Segment Routing Header (SRH) locator with a mapped flexible algorithm, include the `algorithm` statement at the `[edit protocols isis source-packet-routing srv6 locator]` hierarchy level.

To configure a topology-independent loop-free alternate (TI-LFA) backup path for SRv6 in an IS-IS network,include the `transit-srh-insert` statement at the `[edit protocols isis source-packet-routing srv6]` hierarchy level.

[See How to Enable SRv6 Network Programming in IS-IS Networks.]

- **Support for flexible algorithm in OSPFv2 for segment routing traffic engineering (ACX5448, ACX710, MX204, MX104, MX480, MX960, MX10003, MX2020, and PTX10001)**—Starting in Junos OS Release 21.1R1, you can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path based on traffic engineering metric to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

  To define a flexible algorithm, include the `flex-algorithm` statement at the `[edit routing-options]` hierarchy level.

  To configure a device to participate in a flexible algorithm, include the `flex-algorithm` statement at the `[edit protocols ospf source-packet-routing]` hierarchy level.

  [See How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering.]

- **Support for strict SPF and IGP shortcut (ACX710, MX960, MX10008, MX2020, PTX5000, and PTX1000)**—Starting in Junos OS Release 21.1R1, you can configure segment routing algorithm 1 (strict SPF) and advertise its SIDs in IS-IS link-state PDU (LSPDU) and use these SIDs to create SR-TE tunnels to forward the traffic by using the shortest IGP path to reach the tunnel endpoint while avoiding loops. You can also specify a set of prefixes in the import policy, based on which the tunnel can redirect the traffic to a certain destination. You can use algorithm 1 (strict SPF) along with algorithm 0 (default SPF) by default when Source Packet Routing in Networking (SPRING) is enabled.

  [See How to Enable Strict SPF SIDs and IGP Shortcut, prefix-segment, and source-packet-routing.]

Services Applications

- **Support for displaying the timestamp in syslog (MX Series routers with MS-MPC, MS-MIC, and MX-SPC3)**—Starting in Junos OS Release 21.1R1, you can enable system log (syslog) timestamps in local system timestamp format or UTC format.

On routers with MS-MPC, you can override the default UTC timestamp to local system timestamp format by configuring the new statement, `syslog-local-system-timestamp`, at the `edit interfaces` *`ms-interface`*|*`ams-interface`*`services-options` hierarchy level.

On routers with MX-SPC3 cards, you can override the default local system timestamp in syslog to UTC format by configuring the existing statement,`utc-timestamp`, at the `edit interfaces` *`vms-interface`*|*`ams-interface`*`services-options` hierarchy level or at the `[edit services` *`service-set-name`*`syslog` hierarchy level.

For the routers with MX-SPC3 cards, starting in Release 21.1R1 you can configure the `utc-timestamp` statement at the `edit interfaces` *`vms-interface`*|*`ams-interface`*`services-options` hierarchy level. In earlier releases, we support this statement at the `[edit services` *`service-set-name`*`syslog` hierarchy level.

[See syslog (Services Service Set).]

- **Enhancements to DNS sinkhole feature (MX240, MX480, and MX960 routers with MS-MPC and MX-SPC3)**—Starting in Junos OS Release 21.1R1 as part of the DNS sinkhole feature enhancements, you can:
  - Configure new actions for a DNS request for a disallowed domain-alert, accept, drop, and drop-no-log.
  - Configure domain names and actions for multiple tenants such that domain feeds can be managed on a per tenant basis.
  - Configure hierarchical domain feed management per profile, dns-filter-template or dns-filter-term.
  - Exempt domain feeds at the IP, subnet, or CIDR level.

  [See DNS Request Filtering for Disallowed Website Domains.]

- **TWAMP Light IPv4 support (MX Series, PTX Series)**—Starting in Junos OS Release 21.1R1, we support the Two-Way Active Measurement Protocol (TWAMP) Light, as defined in Appendix I of RFC 5357. TWAMP Light is a stateless version of TWAMP, where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away.

  [See twamp.]

- **Support for the any firewall filter family and the Layer 2 firewall filter families for inline monitoring services (MX Series)**—Starting in Junos OS Release 21.1R1, you can configure the `any`, `bridge`, `ccc`, `mpls`, or `vpls` family firewall filter with the term action `inline-monitoring-instance` *`inline-monitoring-instance-name`*.

  [See Inline Monitoring Services Configuration .]

Software-Defined Networking (SDN)

- **Configure SLCs and assign them to GNFs (MX2010 and MX2020)**—Starting in Junos OS Release 21.1R1, in an external server-based Junos node slicing setup, you can additionally configure logical partitions (called sub line cards or SLCs) of the MX2K-MPC11E line card and assign each partition to different guest network functions (GNFs). See Sub Line Card Overview for details. You can create two SLCs on an MX2K-MPC11E. An SLC functions like an independent line card.

  In Junos OS Release 21.1R1, SLCs do not support handling of failures of links between Control Boards and the server, graceful Routing Engine switchover on BSYS and GNF, and unified in-service software upgrade.

  > **NOTE**: In Junos OS Release 21.1R1, Junos node slicing is not multi-version interoperable with previous releases of Junos OS (whether or not SLCs are configured). So, for any GNF in a node-sliced system to run Junos OS Release 21.1R1, all other GNFs and BSYS must also run Junos OS Release 21.1R1.

  [See Configuring Sub Line Cards and Assigning Them to GNFs.]

- **Support for ECMP on multiple flexible routes (MX Series routers with MPC10 and MPC11 line cards)**—Starting in Junos OS Release 21.1R1, MX Series routers with MCP10 or MPC11 line cards support traffic load balancing over multiple flexible routes with 64-way ECMP. A flexible route is a static route with a tunnel encapsulation profile that has the flexible tunnel interface (FTI) attribute. Multiple flexible routes can go through the same logical interface. You can install flexible routes on Juniper gateway devices using Juniper Extension Toolkit (JET) APIs.

  When the router receives a packet with the flexible route as the destination address, it processes the packet using the profile associated with a flexible route, and load-balances the traffic across multiple flexible routes based on the traffic priority.

  Use the `show route` and `show route extensive` CLI commands or the `get-route-information` RPC/NETCONF command to view details about a flexible route for a destination address.

  [See Understanding Programmable Flexible VXLAN Tunnels.]

Software Installation and Upgrade

- **Support for signed third-party application installation (MX10003, MX10008, QFX5210, QFX10002, and QFX10008 routers with VM host architecture)**—Starting in Junos OS Release 21.1R1, you can install signed third-party application installation and carry over the application between upgrades.

  The backup of third-party package occurs during upgrade. Hence, the package is restored even if the installed package is deleted or uninstalled before a reboot. However, as the third party package

restoration depends on the contents saved on the disk during upgrade and the configuration to allow the package to be installed, restoration is not possible when

- Configuration is removed after upgrade
- Content is removed due to deletion by configuring `request vmost zeroize` command

On platforms where jinstall-host.tgz images are installed, the minimum space required for the backup is 250MB. After backup, if the free space available is less than 200MB, the backup would be deleted to make space for upgrade. On platforms where junos-vmhost images are installed, the minimum space required for backup of third party unbundled packages is 1200MB. After the backup, if the free space is less than 512MB, the backup would be deleted to free up space for upgrade.

[See Installing, Upgrading, Backing Up, and Recovery of VM Host.]

- **request system software status command (MX480, MX960, MX2010, MX2020, SRX1500, SRX4100, SRX4400, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, you can use the CLI command `request system software status` to view the status of the software package installation or uninstallation on the local Routing Engine.

Subscriber Management and Services

- **L2TP session lockout support (MX Series)**—Starting in Junos OS Release 21.1R1, you can specify the `lockout-result-code` and `lockout-error-code` options to control the L2TP access concentrator (LAC) behavior in the Layer 2 Tunneling Protocol (L2TP) session lockout state.

  [See lockout-timeout (L2TP Destination Lockout).]

- **Support for PWHT (over EVPN-VPWS, on a transport logical interface) with subscriber management (BNG) service logical interfaces (MX Series routers)**—Starting in Junos OS Release 21.1R1, you can deploy broadband network gateways (BNGs) that are connected to aggregation networks running EVPN-VPWS. You configure pseudowire headend termination (PWHT) on a transport logical interface that is on the pseudowire subscriber interface. The BNG pops the EVPN and VPWS headers and terminates subscribers at Layer 2.

  This feature includes support for:

  - All broadband features available on PWHT on MX Series routers
  - Single-homed EVPN-VPWS with the pseudowire subscriber interface anchored to a logical tunnel (LT) interface
  - Choice of whether or not to use a control word

Virtual Chassis

- **Ephemeral configuration database support for synchronous commit synchronize operations on MX Series Virtual Chassis (MX240, MX480, MX960, MX2010, and MX2020)**—Starting in Junos OS Release 21.1R1, you can configure the ephemeral database to execute commit synchronize operations using a synchronous commit model on MX Series Virtual Chassis. The synchronous commit model enables you to reliably use the ephemeral database on devices that have graceful Routing Engine switchover (GRES) or non-stop routing (NSR) enabled. To use the synchronous commit model for the ephemeral database, configure the `commit-synchronize-model synchronous` statement at the `[edit system configuration-database ephemeral]` hierarchy level.
  [See Understanding the Ephemeral Configuration Database.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for MX Series routers.

What's Changed in Release 21.1R3

IN THIS SECTION

EVPN

- **Output for show Ethernet switching flood extensive**—The output for show ethernet-switching flood extensive now displays the correct next-hop type for Virtual Ethernet and WAN mesh group in an EVPN-VXLAN network as unilist. Previously, the output for show ethernet-switching flood extensive would misidentify the next-hop type as composite.

- **Log messages are removed (MX Series)**—When PTP aggregate Ethernet primary is configured, and PTP Aggregate Ethernet secondary is not configured, the log message **Profiles are being modified** is removed.

General Routing

- **No support for PKI operational mode commands on the Junos Limited version (MX Series routers, PTX Series routers, and SRX Series devices)**—We do not support `request`, `show`, and `lear` PKI-related operational commands on the limited encryption Junos image ("Junos Limited"). If you try to execute PKI operational commands on a limited encryption Junos image, then an appropriate error message is displayed. The `pkid` process does not run on Junos Limited version image. Hence, the limited version does not support any PKI-related operation.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Prior to this PR, the source and destination address pair had to be unique among all the FTI tunnels regardless of the tunnel encapsulation type.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

  [See request system scripts refresh-from and cert-file (Scripts).]

Routing Protocols

- To achieve consistency among resource paths, the resource path /mpls/signalling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter[ip-addr='address']/state/counters[name='name']/out-pkts/ is changed to /mpls/signaling-protocols/segment-routing/aggregate-sid-counters/aggregate-sid-counter[ip-addr='address']/state/counters[name='name']/. The leaf "out-pkts" is removed from the end of the path, and "signalling" is changed to "signaling" (with one "l").

Subscriber Management and Services

- **New output fields for subscriber management statistics (MX Series)**—If you enable the enhanced subscriber management, the non-DHCPv4 bootstrap protocol (BOOTP) requests might not get processed even if you configure the DHCP relay or server with the `overrides bootp-support` statement at the `edit forwarding-options dhcp-relay` hierarchy level. To monitor the DHCP transmit and receive packet counters, we've introduced the following output fields for `show system subscriber-management statistics dhcp extensive` operational command.

  - `BOOTP boot request packets received`
  - `BOOTP boot reply packets received`
  - `BOOTP boot request packets transmitted`
  - `BOOTP boot reply packets transmitted`

  [See show system subscriber-management statistics.]

What's Changed in Release 21.1R2

IN THIS SECTION

General Routing

- **VLAN isolation disabled by default (MX480, MX960, MX2008, MX2010, and MX2020)**—For Junos node slicing, the internal control plane no longer isolates GNFs from each other by default. The internal network has sufficient bandwidth to accommodate GNFs without needing to isolate GNFs from each other. However, if you want to isolate the internal traffic of each GNF from all others, you must configure the `set chassis network-slices vlan-isolation` CLI configuration statement (which is applicable for all uses except with sub line cards) on all the Routing Engines of the BSYS and GNFs and then reboot the chassis. If you want to configure the sub line card feature, you must ensure that VLAN isolation is disabled. We have deprecated the configuration statement `no-vlan-isolation`.

  [See vlan-isolation.]

- **Commit checks against incorrect configuration of SLC values (MX2020 and MX2010)**—We have introduced commit checks against incorrect configuration of sub line cards (SLCs). While configuring SLCs, if you specify any incorrect values (for example, unsupported Packet Forwarding Engine ranges, CPU cores, or DRAM values), the configuration commit fails with an appropriate message to indicate the error.

  [See Configuring Sub Line Cards and Assigning Them to GNFs.]

- **Support for multiple proxy-id list (MX5, MX10, MX40, MX80, MX104, MX240, MX480, MX960, MX2008, MX2010, and MX2020)**—MX Series routers does not support ID list except for the following two cases:

  - MX Series routers accept any-any traffic selector in proxy-id list from the remote device that supports ID lists.

  - MX Series routers accept the ID list if list can be reduced by removing duplicates to specific ID. For example, reduce ID list having 80.0.0.1 and 80.0.0.0/24 to super set ID 80.0.0.0/24.

    list(any:0,ipv4(any:0-65535,[0..3]=80.0.0.1), ipv4_subnet(any:0-65535,[0..7]=80.0.0.0/24))

- **ISSU is not supported**—Unified in-service software upgrade (ISSU) is not supported when clock synchronization is configured for Precision Time Protocol (PTP) and Synchronous Ethernet.

EVPN

- **Support for displaying SVLBNH information**—You can now view shared VXLAN load balancing next hop (SVLBNH) information when you display the VXLAN tunnel endpoint information for a specified ESI and routing instance by using `show ethernet-switching vxlan-tunnel-end-point esi` *`esi-identifier`* `esi-identifier instance` *`instance`* `svlbnh` command.

Interfaces and Chassis

- **Blocking duplicate IP detection in the same routing instance (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, and SRX Series)**—Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. Refer to the table mentioned in the topic inet (interfaces). When you try to configure same IP on two logical interfaces inside same routing instance, the commit will be blocked with the error displayed as shown below:

```
[edit]
user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24

[edit]
user@host# commit
commit complete

[edit]
user@host# set interfaces ge-0/0/2 unit 0 family inet address 2.2.2.2/24

[edit]
user@host# commit
[edit interfaces ge-0/0/2 unit 0 family inet]
  'address 2.2.2.2/24'
     identical local address found on rt_inst [default], intfs [ge-0/0/2.0 and ge-0/0/1.0],
family [inet].
error: configuration check-out failed
```

  [See inet(interfaces).]

Layer 2 Ethernet Services

- **Active leasequery-based bulk leasequery (MX Series)**—The `overrides always-write-option-82` and `relay-option-82 circuit-id` configurations at the `[edit forwarding-options dhcp-relay]` hierarchy level are not mandatory for active leasequery-based bulk leasequery. For earlier releases, the `overrides always-write-option-82` and `circuit-id` configurations are mandatory for active leasequery-based bulk leasequery. For regular bulk leasequery between relay and server without any active leasequery, the `overrides always-write-option-82` and `relay-option-82 circuit-id` configurations are mandatory.

  [See bulk-leasequery (DHCP Relay Agent).]

- **Link selection support for DHCP**—We have introduced the `link-selection` statement at the `[edit forwarding-options dhcp-relay relay-option-82]` hierarchy level, which allows DHCP relay to add suboption

5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope.

Prior to this release, the DHCP relay dropped packets during the renewal DHCP process and the DHCP server used the leaf's address as a destination to acknowledge the DHCP renewal message.

[See relay-option-82.]

Network Management and Monitoring

- **Changes in contextEngineID for SNMPv3 INFORMS (PTX Series, QFX Series, ACX Series, EX Series, MX Series, and SRX Series**—Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  [See SNMP MIBs and Traps Supported by Junos OS.]

- **Enhancement to the snmp mib walk command (PTX Series, QFX Series, EX Series, MX Series, SRX Series)**— The ipv6IfOperStatus field displays the current operational state of the interface. The noIfIdentifier(3) state indicates that no valid Interface Identifier is assigned to the interface. This state usually indicates that the link-local interface address failed Duplicate Address Detection. When you specify the 'Duplicate Address Detected' error flag on the interface, the new value (noIfIdentifier(3)) is displayed. Previously, the snmp mib walk command did not display the new value (noIfIdIdentifier(3)).

- **Change in OID ifHighSpeed**—Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

  [See SNMP MIBs and Traps Supported by Junos OS.]

What's Changed in Release 21.1R1

IN THIS SECTION

VPNs | 147

General Routing

- **Updates to ON-CHANGE and periodic dynamic subscriber interface metadata sensors (MX Series routers and EX9200 line of switches)**—We've made the following updates to the `/junos/system/subscriber-management/dynamic-interfaces/interfaces/meta-data/interface<user-typing>sid='<variable>sid-value</variable>'</user-typing>/` sensor:

  - Notifications are sent when subscribers log in on either IP demux or VLAN demux interfaces. In earlier releases, login notifications are sent only for IP demux logins.

  - The *interface-set* end path has been added to the logical interface metadata. The interface-set field appears in both ON-CHANGE and periodic notifications. In earlier releases, this field is not included in the sensor metadata or notifications.

  [See gRPC Sensors for Subscriber Statistics and Queue Statistics for Dynamic Interfaces and Interface-Sets (Junos Telemetry Interface).]

- **New commit check for MC-LAG (MX Series)**—We've introduced a new commit check to check the values assigned to the redundancy group identification number on the MC-AE interface (`redundancy-group-id`) and ICCP peer (`redundancy-group-id-list`) when you configure multichassis link aggregation groups (MC-LAGs). If the values are different, the system reports a commit check error. In previous releases, if the configured values were different, the l2ald process would crash.

  [See iccp.]

- **Support for unicast ARP request on table entry expiration**—You can configure the device to send a unicast ARP request instead of the default broadcast request when an ARP table entry is about to expire. The retry requests are unicast at intervals of 5 seconds. Without this option, the retry requests are broadcast at intervals of 800 milliseconds. This behavior reduces ARP overall broadcast traffic. It also supports the use case where access nodes are configured not to forward broadcast ARP requests toward customer CPEs for security reasons and instead translate ARP broadcasts to unicast requests. To confirm whether the device is configured, you can issue the following command: `show configuration system arp | grep unicast-mode-on-expire`

  [See arp.]

- **Update to the show chassis errors active output (MX2010 and MX2020 routers with MPC11E)**—We have updated the `show chassis errors active` output for the MPC11E line card (MX2K-MPC11E) to display the correct error information. Previously, this CLI command displayed duplicate or incorrect output when the MPC11E line card is not installed in slot 0 of the MX2010 or MX2020 routers.

[See show chassis errors active.]

- **CLI commit error resolved**—CLI commit error for configuration statement `no-filter-check` for `family any` port mirroring is now resolved.

  [See no-filter-check.]

Interfaces and Chassis

- **Hardware-assisted timestamping**—By default, hardware assistance is used for timestamping Ethernet frame delay frames on AFT-based MX Series line cards, even if `hardware-assisted-timestamping` is not configured.

  [See Enabling the Hardware-Assisted Timestamping Option.]

- **Change in <range> XML tag (MX480)**—We've changed the `<range> string </range>` XML tag to `<transport-range> <transport-range-info> string </transport-range-info> <transport-range-suspect-flag> string </transport-range-suspect-flag> <transport-range-reason> string </transport-range-reason> </transport-range>` under the [`show interfaces transport pm optics current <interface> | display`] hierarchy in the XML output. Hence, the new XML tags that associate the values to the range-info, range-suspect-flag, and range-reason tags map the information to the given `show interfaces transport pm optics current interface | display` entry.

  [See Supported OTN Options on MX Series Routers.]

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

[See invoke() Function (SLAX and XSLT).]

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Layer 2 Ethernet Services

Active leasequery-based bulk leasequery (MX Series)—The `overrides always-write-option-82` and `relay-option-82 circuit-id` configuration at the `[edit forwarding-options dhcp-relay]` hierarchy level is not mandatory for active leasequery-based bulk leasequery. In releases before Junos OS Release 21.1R1, the `overrides always-write-option-82` and `circuit-id` configurations are mandatory for active leasequery-based bulk leasequery. For regular bulk leasequery between relay and server without any active leasequery, the `overrides always-write-option-82` and `relay-option-82 circuit-id` configurations are mandatory.

[See bulk-leasequery (DHCP Relay Agent).]

Network Management and Monitoring

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the `[edit system services netconf hello-message yang-module-capabilities]` hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the `[edit system services netconf netconf-monitoring netconf-state-schemas]` hierarchy level.

  [See hello-message and netconf-monitoring.]

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

User Interface and Configuration

Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)—The Junos OS CLI exposes the `verbose` statement at the `[edit system export-format json]` hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from `verbose` to `ietf` starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `[edit system export-format json]` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

[See export-format.]

VPNs

View the traffic selector type for an IPsec tunnel (SRX Series and MX Series)—You can run the `show security ipsec security-associations detail` command to display the traffic selector type for a VPN. The command displays `proxy-id` or `traffic-selector` as a value for the `TS Type` output field based on your configuration.

[See show security ipsec security-associations.]

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for MX Series routers.

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

General Routing

- When the device is up and running for a long time, there is a possibility FS gets bad blocks and it is accumulated. When any change is done to it, it reloads and tries to recover the bad blocks from the FS. PR910445
- On the MX104 router, scheduler slip is observed when configuration changes are committed. PR1361250
- Traffic stops after the volume limit is reached, but the traffic resumes after an aggregated Packet Forwarding Engine interface fails. PR1463723
- For two tunnels with same TS, all traffic is forwarded on the tunnel that comes up later. PR1467364
- LFM might flap during MX Virtual Chassis unified ISSU to and from this release. PR1516744
- BGP neighbor flaps during the primary NG-RE reboot, although GRES or NSR is enabled. PR1524791
- When an image with the third-party SDK upgrade (6.5.x) is installed, the CPU utilization might go up by around 5 percent. PR1534234
- The NPC process continuously generates core files at `Trinity_Ktree::Trinity_FourWayBlock`, `Trinity_Ktree::walkSubTree` due to the next-hop memory exhaustion with the next-hop explosion. The rpd and srrd processes start hogging and the system becomes unstable. PR1538029
- Guidelines are needed for enabling or restarting fib-streaming in low-memory conditions. PR1540478
- Issuing the `help apropos` command in configuration mode is going to cause an mgd core. The mgd process will come up and as long as the command is not issued again, the core will not occur. PR1552191
- Unified ISSU is not supported. There is a major SDK upgrade from 6.3.2 to 6.5.16, due to which the warm boot feature needed for unified ISSU is not supported by vendor. PR1554915
- On the MX10003 routers, after a subscriber logs in, the subscriber is allowed to delete the entire AGF service stanza and commit the configurations, leaving the subscribers stuck in the **Active** state permanently. The user must not be allowed to delete the AGF Service configurations with active UEs. PR1555031

- With the IPIP tunnel feature, the `show dynamic-tunnels database statistics` command output shows extra packet counts (that is, sampled packets when sFlow is enabled). PR1555922
- sFlow egress sampling of MPLS packets is not supported on MX Series platforms. PR1556659
- `Resource deadlock avoided` messages are observed during software add. There is no functionality impact. PR1557468
- On the MPC11E line cards, the following error message is observed:

  ```
  ppman - PPM:RPC - Error message <url
  ```

  PR1559434
- On the MPC11E line cards, the following error message is observed:

  ```
  l2tp-sfd[11852]: [Error] L2TP-SFD & CFMMAN & VBFMAN & RPC-SERVICE
  ```

  PR1559440
- The rpd process generates core file if the `use-for-shortcut` command is configured on an SR-TE tunnel that uses an SR Algo 0 prefix SID. PR1578994
- SyncE to PTP noise transfer passes for 400 ns p-p amplitude and frequency of 1 Hz but fails for 200 ns p-p amplitude and frequency of 0.005 Hz.PR1566291
- G.8273.2 transient response test fails. The issue exists for legacy line cards also. PR1566354
- With T-BC across multiple line card, average time error (cTE) test fails as there are other delays introduced causing phase variation across line cards. PR1567662
- On deactivating aggregated Ethernet interface, we might see a traffic loss of greater than 2 seconds. This behavior is seen due to order in which the messages are processed in FPC wherein next hop change/delete is processed prior to interface down event. PR1614508
- Spike in the rpd usage is expected because of the very large scale and OCST in general, but it does not affect any rpd functionality as telemetry streaming is the least priority task in the rpd. PR1614978

Infrastructure

- Junos OS Release 21.1 and earlier releases run FreeBSD version 11 whereas from Junos OS Release 21.2 onward, the FreeBSD version is 12. Software upgrade to 21.2 or later releases from 21.1 or earlier releases must use `no-validate` CLI option during software image upgrade process. PR1586481
- During high route churn in scaled configurations, peer buffer becomes full and statistics request might get dropped by kernel because of their low priority as route updates are prioritized by kernel. This is a design limitation to prioritize route updates over statistics. The application is expected to retry in this scenario after route churn settled down. PR1607362

Interfaces and Chassis

- For MC-LAG to work properly, the mc-ae interface should be configured on both the PE devices. A scenario where the mc-ae interface is deleted, deactivated, or not configured on one of the devices is a case of misconfiguration. Juniper Networks does not support such a scenario because it can lead to traffic loss and other unexpected behavior. PR1536831
- On the MPC10 line cards, DMRs or SLRs are not received with an EVPN up MEP on the aggregated Ethernet interface with normalization. PR1543641
- Packet loss is seen in the scaled setup with 296 LM sessions with iterator cycle time interval (100 ms). It seems there is degradation in scale number (OAM packet rate at ~5500). At this qualified PPS, now LMR packet loss is observed, but the functionality seems to be fine. To avoid LMR packet loss, reduce the scale number, and OAM packet rate should be less than 5500 PPS. PR1561397

MPLS

- The rpd process might crash. PR1461468

Network Management and Monitoring

- SNMP link up trap message is not observed after a line card reboots when scaled interfaces are present. PR1507780
- The `set system no-hidden-commands` configuration blocks NETCONF sessions. As a workaround, customer can disable with the `no-hidden-commands` statement. PR1590350

Platform and Infrastructure

- On MX Series platforms, under an EVPN environment, packets routed using IRB interface might not be fragmented due to media maximum transmission unit (MTU) problem. PR1522896
- Traffic convergence is more than 50ms after disabling core-facing links on primary PE device when the core links on egress PE device are on MPC11 line card. PR1562761

Routing Protocol

- Convergence time is high when the igmp snooping configuration is deleted. PR1550523
- With maximum number of logical interfaces (4000 GRE tunnel per Packet Forwarding Engine) with the following configuration:

  1) family inet, associated source, and destination for each tunnel.

  2) Configure allow-fragmentation statement on one endpoint of the tunnel and configure reassemble-packets on the other endpoint of the tunnel.

  If we do `deactivate chassis fpc <slot>`, we might see SLIP messages while testing inline GRE reassembly feature with GRE interface scaling for MX Series. PR1581042

User Interface and Configuration

- Unsupported options are displayed under the `restart` commands. PR1545558

VPNs

- In some scenario (for example, configuring firewall filter), sometimes device might show obsolete IPsec SA and NHTB entry even when the peer tear down the tunnel. PR1432925
- Traffic loss is observed on starting the traffic with PIM disabled under the ingress primary UMH. PR1562759

Open Issues

IN THIS SECTION

Learn about open issues in this release for MX Series routers.

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

General Routing

- Some non-fatal interrupts (for example, CM cache or AQD interrupts) are logged as fatal interrupts. The following log messages will be shown on CM parity interrupt:

```
fpc0 TQCHIP 0: CM parity Fatal interrupt,Interrupt status:0x10
fpc0 CMSNG: Fatal ASIC error, chip TQ
fpc0 TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180010 msecs
TQCHIP 0: CM cache parity Fatal interrupt has occurred 181 time(s) in 180005 msecs
```

  PR1089955

- On MX104 platforms, when using the `snmpbulkget` or `snmpbulkwalk` (for example, used by the SNMP server) on a chassisd-related component (for example, jnxOperatingEntry), high CPU usage for chassis process and slow response might be seen because of a hardware limitation, which might also lead to a query time out on the SNMP client. In addition, the issue might not be seen while using an SNMP query for interface statistics. As a workaround, to avoid the issue, use either of the following approaches:

  Use snmpget or snmpwalk instead of snmpbulkget or snmpbulkwalk and include the -t 30 option when doing the SNMP query. For example, `snmpget -v2c -c XX -t 30`.

  Use the -t 30 option with snmpbulkget or snmpbulkwalk. For example, `snmpbulkget -v2c -c XX -t 30`.
  PR1103870

- On MX platforms with FPC-PTX-P1-A or FPC2-PTX-P1A, you might encounter a single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message gets reported: `Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0 Fatal Errors - TQ Chip Error code: 0x50002 Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error code: 0x50002`.

  The Junos OS chassis management error handling detects such a condition, raises an alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper Networks support representative if the issue persists even after the FPC restart. PR1254415

- If a vmhost snapshot is taken on an alternate disk and there is no further vmhost software image upgrade, the expectation is that if the current vmhost image gets corrupted, the system boots with the alternate disk so the user can recover the primary disk to restore the state. However, the host root file system and the node boots with the previous vmhost software instead of the alternate disk. PR1281554

- When you issue a `show interface` command to check the interface details, the system does 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

- With aggregated Ethernet bundle or ECMP next hops configured with the adaptive load balancing feature requests a large chunk of jnh counter memory . If allocation requests are spread over an interval of time, then the memory allocator might not be able to handle all these requests and error messages are reported. There is no impact on traffic. PR1329704

- Source MAC and TTL values are not updated for routed multicast packets in an EVPN VXLAN scenario. PR1346894

- The backup Routing Engine might crash after GRES occurs continuously for more than 10 times. PR1348806

- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. PR1362609

- The following log message might be seen on FPC with WINTEC mSATA SSD:

  ```
  SMART ATA Error Log Structure error: invalid SMART checksum
  ```

  PR1354070

- A few xe- interfaces go down with the following error message: `if_msg_ifd_cmd_tlv_decode ifd xe-0/0/0 #190 down with ASIC Error`. PR1377840

- The ping command might show variable latency values. This is expected for host generated ICMP traffic due to the design of the Packet Forwarding Engine queue polling the packets from ASIC. PR1380145

- Due to a transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possible unnecessary hardware replacements. This change applies to all platforms using Hybrid Memory Controller (HMC). PR1384435

- Modifying the underlying interface on a demux0 interface with subscribers present on the underlying interface causes the FPC to generate core files. In the procedure to edit underlying-interface on a demux0, do the following check:

  Verify that there are no subscribers existing on the underlying interface configured on the demux0.

  Subscribers need to be moved out of the underlying interface before editing the underlying interface under demux0. PR1396157

- The PTP master and the PTP slave port configuration accept only the PTP packets with multicast MAC address according to the port settings. If forwardable multicast is configured, only PTP packets with the forwardable MAC address is accepted and the non-forwardable is dropped. Similarly, if the link-local multicast is configured, only the PTP packets with the non-forwardable MAC address is accepted and forwardable is dropped. PR1442055

- On MX10003 routers with Virtual Chassis, access facing FPC's CPU stays at 100 percent for 5 to 6 minutes after a configuration change. PR1447003

- On VXLAN VNI (multicast learning) scaling, traffic issue is seen from VXLAN tunnel to Layer 2 interface. PR1462548

- In DNS filtering, when DNS requests are sent from the server and implicit filters as well as routes to the service PIC are configured, it causes the DNS packets to loop. As a workaround, configure either static routes or implicit filters for forwarding DNS traffic to service PIC. It solves DNS packet looping issue. PR1468398

- On MX Series platforms with the 3D 20x 1GE MIC installed, after performing ISSU, the FPC equipped with the MIC might crash and interfaces stay down. Due to this issue, the traffic on the MIC will be impacted. PR1480212

- On MX204 and MX10003 routers with the MPC7E, MPC8E, MPC9E, MPC10E, and JNP10K-LC2101 line cards, the following syslog error appears occasionally: `unable to set line-side lane config (err 30)`. This does not impact the service and can be ignored. PR1492162

- When the `show pfe filter hw filter-name` *`filter name`* command is issued, the command fails to retrieve the Packet Forwarding Engine programming details of the filter. PR1495712

- After the backup Routing Engine halts, CB1 goes offline and comes back online. This leads to rebooting of the backup Routing Engine and it shows the reboot reason as `0x1:power cycle/failure`. There is no other functional impact due to this issue. PR1497592

- When a VLAN member is specified as a string, the `IF_MSG_IFL_VADDR` TLV is not generated with the VLAN infoformation, and the MX Series with MPCs or MICs afttriostream is not updated with the nativevlanId and nativevlanenable flags. Thus, the packet is still treated as untagged, and when it reaches the trunk egress interface, it is dropped because the trunk interface does not allow untagged traffic to pass through. The issue is specific to platforms with ZT line cards. As a workaround, configure the `interface-vlan-members` statement with only numeral value only for VLANs. The VLAN members with input as a string is not supported in this release. PR1506403

- A 10-Gigabit Ethernet interface configured with WAN-PHY framing might flap continuously if the hold-down timer is set to 0 (which is the default). This is not applicable to an interface with the default framing LAN-PHY. PR1508794

- On a fully scaled system where all the slices are utilized by different families of CLI filters, if we try to delete one family and add or change another family with a higher number of filter terms, which

requires either expansion of the filter or creation of a new filter, the Packet Forwarding Engine fails to add the new filter as we are getting out of sequence messages. The add or change of the filter is called earlier than the delete of another filter will free up the slices. PR1512242

- A 35 seconds delay is added in reboot time. PR1514364
- When an AMS physical interface is configured for the first time or any member of the AMS bundle is removed or added, the PICs on which the members of AMS bundle are present go for a reboot. There is a timer running in the AMS kernel which is used as a delay for the PIC reboot to complete and once that timer expires, AMS assumes that the PICs might have been rebooted, and it moves into next step of AMS finite state machine (FSM). In a scaled scenario, this rebooting of the PIC is delayed due to DCD. This is because when a PIC goes down, DCD is supposed to delete the physical interfaces on that PIC and the PIC reboot happens. But DCD is busy processing the scaled configuration and the physical interface deletion is delayed. This delay is much greater than the timer running in AMS kernel. When the timer expires, the FSM in AMS kernel incorrectly assumes the PIC reboot would be completed by then, but the reboot is still pending. By the time DCD deletes this physical interface, the AMS bundles are already up. Because of this, there is a momentary flap of the bundles. PR1521929
- The rpd sensors generate core file during defer-continue case on a network churn. This will be a timing issue and will happen only when a particular node sensor information is being rendered and the same node went through some modification. PR1526503
- On the MX Series platforms with next generation Routing Engine installed, after upgrading the Intel i40e-NVM firmware to version 6.01, the FRUs disconnection alarms might be seen along with traffic loss. Refer to the TSB17603 to upgrade Junos OS software and Intel i40e-NVM firmware. PR1529710
- On MX150 routers, the following error messages are seen in the messages log file for the interfaces that have SFP installed in them: `fpc0 FAILED(-1) read of SFP eeprom for port: 13`. PR1529939
- FIPS mode is not supported. PR1530951
- Ping command does not work even though the ARP entry is present during continuous script executions due to BRCM KBP issue. PR1533513
- After performing a unified ISSU in a Junos OS node slicing, the unified ISSU unsupported field replaceable unit (FRU) will stay offline until it brings back to online manually once ISSU is finished. This issue causes a service or traffic impact for the offline FRUs. PR1534225
- When an image with the third party SDK upgrade (6.5.x) is installed, the CPU utilization might go up by around 5 percent. PR1534234
- Flapping might be observed on channelized ports of MX Series routers during ZTP, when one of the ports is disabled on the supporting device. PR1534614

- In rare instances, when the et- interface gets stuck and remains down between two particular ports on MPC5E, MPC4E, and CXP MIC line cards. The MAC chip in the line card goes to a condition where the EDC convergence state (adaptive algorithm state machine error) in the firmware remains at `tracking` while the et- interface is stuck and remains down. PR1535078

- The `request system software validate` command is disabled currently for Junos OS Release 19.4 and later. You can validate the same using the `request system software add` command. PR1537729

- The `Socket to sflowd closed` error comes up when the ukern socket to sflowd daemon (server) is closed. The error is rectified by itself as the client successfully reestablishes the connection in the subsequent attempts. When these errors are consistent, it indicates a communication issue between sflowd and the sFlow running on the FPC. PR1538863

- In an EVPN-VXLAN scenario with the Layer 2 and Layer 3 multicast configurations, the vmcore process generates the core file on the primary and backup Routing Engines. PR1539259

- On a scaled MX2020 router with vrf localisation enabled, when 4 million next hop scale and 800000 route scale are available, FPCs might go offline on GRES. Post GRES, router continues to report many fabric related CM_ALARMs. The FPC might continue to reboot and does not come online. Rebooting the primary and backup Routing Engine will help to recover and get router back into stable state. PR1539305

- The following error message occurs when you reboot the device with the enterprise base configurations: `Error BCMX: Failed to add lport 0x0 (unit , port ). -8: Entry exists`. PR1541159

- PTP to PTP noise transfer is passing for impairments profile `400nsp-p_1Hz`, but failing for profile `400nsp-p_0.1Hz` and lower bandwidth profiles as well. The issue is common to 10G also. PR1543982

- After performing upgrade or downgrade on VM host platform, during restarting with the new image, the Wind River Linux (WRL) kernel might go into a deadlock state due to a race condition in advanced configuration and power interface component architecture (ACPICA) module in Linux kernel. This issue might cause the system to get stuck in continuous crashing state. It is a rare timing issue and currently only seen on WRL6 kernel based image during upgrade or downgrade. PR1544875

- Intermittent license check core files are generated during the device initialization. License daemon will restart and start providing the required support. There is no service impact. PR1545175

- A new alarm `network-service mode mismatch between configuration and kernel setting` is introduced. When unified ISSU or normal code upgrade is performed from images without new alarm commit to images with new alarm commit, then the transient false alarm will be seen. PR1546002

- Hardware performance counters might not be correctly exported to the CLI when Packet Forwarding Engines are disabled. This is purely a display issue. PR1547890

- 100G AOC from third-party does not come up after multiple reboots. It recovers after enabling and disabling the interface. PR1548525
- The following error message is observed: `Feb 27 20:26:40 xolo fpc3 Cannot scan phys_mem_size.out. Please collect /var/log/*.out (0;0xdd3f6ea0;-1) (posix_interface_get_ram_size_info): Unknown error: -1`. This log is harmless. PR1548677
- In synce configuration, ESMC transmit is configured or if the `chassis synchronization source` configuration is deactivated or there are no active `chassis synchronization source` configurations present, it might lead to a commit error `esmc-transmit`. To avoid the error, include the `chassis synchronization source`. PR1549051
- On MX10008 and MX10016 platforms, the keepalive value of chassisd socket between chassisd and line card is small. Due to this, when issues like short link-flaps/connection problem occur, the FPC reboots instead of reconnecting, which causes service impact. PR1550917
- After a system reboot, BFD session status is in Init state. It is seen when we have both CFM and BFD configuration on the system and endpoint overlaps between CFM and BFD. PR1552235
- Phone home supports captive portal with factory default configuration. Captive portal is used to enter activation code and to monitor bootstrap status of device using phone home feature. Starting Junos OS Release 20.4, support for captive portal for phone home bootstrap process is removed. PR1555112
- 5M DAC connected between QFX10002-60C and MX2010 does not link up. But with 1M and 3M DAC, this interoperation works as expected. There seems to be a certain SI or link-level configuration on both QFX10002-60C and MX2010. PR1555955
- On the MPC11E line cards in BSYS, commit goes through when unified ISSU is initiated in the GNF.PR1556544
- On the MPC9E line card, core files are generated when SFB becomes online after unified ISSU of a GNF. PR1556627
- On high availability systems, when FPC0 (when node0 is primary) or FPC7 (when node1 is primary) is restarted (for example, with the `request chassis fpc slot <> restart node local` CLI command or due to dcpfe core files on the primary), that might cause FPC1 or FPC8 to restart, which might cause the preexisting TCP sessions to break and might not get reestablished by itself. The TCP sessions might need to be manually reestablished. PR1557607
- On the MX10008 routers, the GRE keepalive adjacency state is down even though the GRE tunnel is in the up state. PR1559200
- VE and CE mesh groups are default mesh groups created for a given routing instance. On adding VLAN or bridge domain, flood tokens and routes are created for both VE and CE mesh-group and flood-group. Ideally, VE mesh-group does not require a CE router where IGMP is enabled on CE

interfaces. MX Series based CE boxes have unlimited capacity of tokens, so this would not be a major issue. PR1560588

- In an MVPN scenario, if the next hop index of a group is not same between primary and backup after a NSR switchover, we might see a packet loss of 250 to 400 ms. PR1561287
- The timingd-lc errors `CdaExprClient: grpc api call ExprServerInfoGet failed" and "CdaExprClient: Failed to fetch server info error:5` are seen on all FPCs after restarting router or FPC. PR1561362
- Due to a race condition, the `show multicast route extensive instance instance-name` output can display the session status as Invalid. Such an output is a cosmetic defect and not an indicative of a functional issue. PR1562387
- Configure the interface hold time to avoid the additional interface flap. PR1562857
- In a rare scenario, SPMB does not reply during FPC online which is moved from SLC mode to full line card mode. The FPC gets stuck as the training is not complete. PR1563050
- When SLC is reconfigured from asymmetric mode to symmetric mode in a single commit, it is possible that on some occasions, one of the SLC shows chassis connection as dropped state. The SLC will come online and no functional impact is seen. PR1564233
- When a MPLS p2mp template is configured over the default_p2mp template, the configuration change does not take effect and the old configuration remains active. PR1564795
- Starting in Junos OS Release 21.1R1, Junos OS will be shipping with python3 (python2 is no longer supported). In ZTP process, if a python script is being downloaded, ensure the python script follows python3 syntax (there are certain changes between python2 and python3 syntax). Also, so far (that is, until Junos OS 20.4R1), the python script had #!/usr/bin/python as the first line (that is, the path of the python interpreter). The same needs to be changed to #!/usr/bin/python3 from Junos OS Release 21.1R1. PR1565069
- In a dual CPE scenario, after RG0 failover, the best path link status shows as PARTIAL SLA VIOLATED instead of SLA MET due to active probe result is incorrect in certain scenarios. PR1565777
- SyncE to PTP noise transfer passes for 400 ns p-p amplitude and frequency of 1 Hz, but fails for 200 ns p-p amplitude and frequency of 0.005 Hz. PR1566291
- G.8273.2 transient response test fails. Issue exists for legacy line cards also. PR1566354
- The chassisd logs flood with the `pic_create_ifname: 0/0/0 pic type F050 not supported` messages for every connected port. The flooding might happen every few seconds. PR1566440
- If the inline services and services are applied to sub line cards (SLCs), some issues might happen during processing these services along with firewall process (dfwd) filter actions. Then it might cause SLCs to reboot and aftd to crash. PR1567313

- With T-BC across multiple line card, average time error (cTE) test fails as there are other delays introduced, causing phase variation across line cards. PR1567662
- Fusion cascade ports must not be hosted on the VPN core facing FPC. When VPN localisation is enabled in fusion or v44 setup, ensure cascade ports (satellite devices) are not part of VPN core facing FPC. PR1567850
- The problem is with L1 node not reflecting correct bandwidth configured for tunnel services. When baseline has 1 G configuration on some FPC or PIC in groups global chassis and if we override with local chassis tunnel service in 10 G bandwidth scaled scenario. Out of 10 Gbps bandwidth configured, only 1 Gbps is allowed per 1 G speed configured in baseline configuration. PR1568414
- Traffic might be dropped on MX Series platforms when the default route is changed in the inet.0 table. It might take 2 to 3 seconds to update in Packet Forwarding Engine . This issue will be recovered automatically. PR1568944
- The PTP clock might fail to be locking and stuck in acquiring state at clock servo. PR1570310
- BUM traffic replication over VTEP is sending out more packets than expected and there seems to be a loop. PR1570689
- Part of the output of the `show ptp lock-status detail` command is missing while changing the interface configuration from the encapsulation Ethernet to the family inet. This issue is not seen every time and issue exists for legacy line cards also. PR1572047
- On all Junos platforms, traffic loss might be observed due to a rare timing issue when performing frequent Interface Bridge Domain (IFBD) configuration modifications. This behavior is seen when the Packet Forwarding Engine receives out-of-order IFBD(s) from Routing Engine and might lead to the fxpc process crash and traffic drop. PR1572305
- When trying to configure a separate rib-group for PIM in VRF, after performing the `commit check`, the following error might be seen: `PIM: ribgroup vrf-mcast-v4 not usable in this context; all RIBs are not in instance vrf`. PR1574497
- When the scheduler configuration is not applied to all 8 egress queues of an interface and one or more egress queues is having buffer size remainder configuration, the distribution of buffer to egress queues with buffer size remainder is not distributed correctly, which might lead to unexpected tail drops. PR1575798
- An alarm is raised due to a transient hardware problem with MIC does not get cleared automatically after MIC restart. PR1576370
- Max ports used is not getting displayed properly for the `show services nat pool pool-name detail` command. PR1576398

- On MX Series platforms with the MPC7E, MPC10E, MX-SPC3, and LC2103 line cards might become offline resulting in complete loss of traffic when the device is running on FIPS mode. The `show chassis fpc pic-status` command can be used to check the status of the line cards. PR1576577
- When a firewall is configured with both discard and port-mirror as actions in the same term, mirrored packets are corrupted. PR1576914
- On MX Series platforms, in a subscriber scenario with scaled around 32,000 connections, the replication daemon might generate core files or stop running, which results in failure on subscriber services on the new Routing Engine after the upgrade or GRES. PR1577085
- When a sub line card (SLC) assigned to a GNF in a node sliced setup generates some PCIe alarms during boot up. This alarm does not have any functional impact and will resolve once the SLC is online. PR1578187
- Snapshot banner message displays to reboot the system from primary disk using the `request node reboot re disk1` command, but the correct command is `request node reboot re0 disk1`. PR1578556
- This issue is caused by /8 pool with block size as 1. When the configuration is committed, the block creation utilizes more memory causing NAT pool memory shortage, which is currently being notified to the customer with syslog tagged RT_NAT_POOL_MEMORY_SHORTAGE. PR1579627
- When MPC11E is sliced into Sub Line Cards (SLC) in a node sliced environment, it is possible that in some instances the multiple times restart of one SLC might cause the complete FPC to restart. This could cause a traffic impact. PR1581107
- When a large number of subscribers attempt to subscribe na-grpcd, core file might be seen. All telemetry subscription connections will close and collectors have to subscribe again. PR1583161
- On MX10003 router, PEM capacity might be incorrectly shown by the `show chassis power` command after a PEM swap. PR1587694
- As part of filter configuration, the out-of-order scenario corner case validation is not handled at Packet Forwarding Engine. Because of this, the aftd process crashes at `dfw_term_dictionary_get_next (term=0x0, dfw=0x7f1431da34c0) at ../../../../../src/pfe/common/applications/dfw/dfw_term.c:1460`. PR1589619
- Minor transient traffic drop will be seen during MBB of RSVP LSP without the `optimize-adaptive-teardown` statement. PR1590656
- The subscriber logical interface might be in a stuck state after the extensible subscriber services manager (ESSM)is deleted. PR1591603
- On all MX Series platforms, changing configuration AMS 1:1 warm-standby to load-balance or deterministic NAT might generate vmcore file and traffic loss might be seen. PR1597386
- Read write lock is not acquired during the sysctl invocation. The assert triggered in the interface state function call leads to go Routing Engine 1 to debug (db>) prompt. PR1598814

- After performing an upgrade, the peer device is rebooted or the peer interface is disabled or enabled, then the SFP-T port might remain in up state but might not forward traffic. PR1600291

- On all Junos platforms, in configurations where a large number of tag next-hops have neighbor discovery (ND6) next hop as underlying next hop, upon refresh of ND6 entry because of any reason, a large number of updates are sent to the Packet Forwarding Engine. This update processing causes a spike in the CPU usage which might hamper some scheduled tasks if they occur simultaneously. PR1600318

- When PTP is on default profile and PTPoE is configured in stateful with ordinary clock-mode configuration is not supported. The below unsupported configuration does not throw commit error.

```
user@host# show protocols ptp
clock-mode ordinary;
stateful {
    interface xe-0/0/0.0 {
        multicast-mode {
            transport {
                ieee-802.3;
            }
        }
    }
}
```

  The stateful port configuration for PTP over Ethernet and default profile is supported only on boundary clock mode and not on ordinary clock mode. As a workaround, change the clock-mode or to remove stateful configuration. PR1601843

- When the interface transitions from down to up, the carrier transition counter value of a particular interface might be incorrect when the peer interface takes longer time to come up. Configuring the `hold-time` for up and down helps to resolve this issue. PR1601946

- Core files will be observed in SPC3 when you change dslite configuration multiple times under service-set. PR1601977

- When static routes are added with gr- interface names, there might be replication issues with MPLS next hops causing backup to generate core files. PR1601996

- On MX Series platforms with MPC10E line card, output bps is not in the expected range on aggregated Ethernet interface for egress traffic. PR1602307

- The convergence time degradation is seen in IS-ISv6, OSPFv2, and OSPFv3 when comparing convergence time with Junos OS Release 21.1R1.5. As it is a convergence time issue, many

components are involved and hence need investigation of rpd, kernel, and Packet Forwarding Engine. PR1602334

- In chassis with mix of MPC10 or MPC11 and MPC1 to MPC9 line cards, and aggregated Ethernet bundle configuration with member links on both MPC10/MPC11 and MPC1 to MPC9, packet loss might be seen for unicast packets on link flap using ifconfig down/up command in Routing Engine shell. PR1604073

- In chassis with mix of MPC10 or MPC11 and MPC1 to MPC9 line cards, and aggregated Ethernet bundle configuration with member links on both MPC10/MPC11 and MPC1 to MPC9, packet loss may be seen for unicast packets on link flap when deleting aggregated Ethernet bundle and adding it again. PR1604450

- In chassis with mix of MPC10 or MPC11 and MPC1 to MPC9 line cards, and aggregated Ethernet bundle configuration with member links on both MPC10/MPC11 and MPC1-MPC9, packet loss might be seen for unicast packets on link flap when deactivate and activate bundle. PR1604800

- When performing downgrade on VM host platform, the following harmless error messages might be seen when issuing the `request vmhost software add` command: `mkdir: cannot create directory '/tmp/partdisk-V6pHko/jrootfs/junos': File exists mkdir: cannot create directory '/tmp/partdisk-V6pHko/jrootfs/vm': File exists mkdir: cannot create directory '/tmp/partdisk-V6pHko/jrootfs/spare': File exists`. PR1605915

- On MX240, MX480, and MX960 routers with both MPC10E line card and MPC2, MPC3, MPC4, MPC5, and MPC6 based FPCs, when the MPC10E line card sends high traffic to MPC4E or other mentioned cards as the destination, the destination line card will not be able to cope up with MPC10E line card traffic flow. PR1606296

- The dfwd core files are generated when accessing ephemeral date base files which is deleted through script. PR1609201

- On all Junos OS platforms, when disabling the physical interface where GRE tunnels is established and performing a GRES. After GRES, enabling the physical interface will cause the BFD to become stuck in init state. PR1609630

- In MX240, MX480, and MX960 routers with SCBE3-MX and enhanced midplane, in some rare cases, if huge traffic is flooded from MPC7, MPC8, and MPC9 line card to MPC2E, MPC3E, MPC4E, and MPC5E line card, and flaps the interface on MPC2E, MPC3E, MPC4E, and MPC5E, it will cause the unexpected request time errors on MPC7, MPC8, and MPC9 line cards because the MPC2E, MPC3E, MPC4E, and MPC5E line cards might not be able to handle such high volume of requests. It causes the Packet Forwarding Engine destinations to become unreachable even when the fabrics are online. Then the Packet Forwarding Engine, SIB, SCBE, FPCs might reboot automatically while these accumulated fabric errors hit the fabric connectivity restoration conditions of the fabric healing process (FHP). PR1612957

- In some NAPT44 and NAT64 scenarios, duplicate SESSION_CLOSE syslog error will be seen. PR1614358

- In Junos subscriber management environment, when the subscriber with the input service-filter configured in the service under dynamic profile fails when modified using Change-of-Authorization (CoA). CoA NAK is received when the input-service-filter is modified. PR1614903
- On all MX Series platforms with MPC10 line card (AFT based MPC), if the filter is created with the resolved filter and deactivating filter attached to interface after MPC reboot, no filter found error might be seen when the device have multiple filters configured across different families. Due to this, filter might not be effective and counter fetch might not work. PR1616067
- The transit IPv4-over-IPv6 encapsulated packets cannot pass through using IP over IP interface. This behavior has been seen on transit packets only. PR1618391
- On all MX platforms running Junos OS with enhanced subscriber management environment, agent circuit identifier (ACI)-based dynamic VLAN session might fail when the size of the PPPoE vendor specific (VS) tags containing ACI and access line characteristics TLVs, exceeds 80 bytes, which would not allow the client to login. PR1619122
- On MX platforms with subscriber management redundancy, if DHCP active lease query (ALQ) is configured without the topology discover and the `no-advertise-routes-on-backup` statement is configured, DHCP ALQ connection might not be established and DHCP subscribers might not be synchronized to backup Broadband Network Gateway (BNG). PR1620544
- When the PHY-sync state of a line card moves to FALSE permanently, it fails to send a degraded clock class to its downstream neighbors. PR1622108
- In a virtual router environment, there is a variance with respect to the traffic being passed through the interface where the accounting is enabled. PR1622514
- Port speed shows as 100G even though chassis configuration is set for 40G. This is just a cosmetic display issue. PR1623237
- This is a product limitation for MX SPC3 with new junos-ike architecture. The issue is seen when we have any-any TS configured and any-any TS negotiated (both in IPv4 and IPv6). As a workaround, do not configure any-any TS when it is sure that negotiated traffic selector for the IPsec tunnel will also be any-any. When there is no TS configured, the scenario might be treated as proxy-id case and bypasses the issue without having any impact on the described scenario. PR1624381
- If option 80 ahead of option 82 in the client's DHCP discover packet, the `auto-configure` feature can not extract the subscriber's agent circuit-ID (ACI) and agent remote-ID (ARI). This leads to authentication failure when creating the dynamic VLAN interface where option 82 is requested. PR1626558
- On all Junos OS platforms, the line card might crash and reload in an EVPN-MPLS scenario when there is a MAC move from local to remote and the request to delete MAC entry is received from remote. Core files are generated and complete traffic loss might be observed until the line card is reloaded. PR1627617

- The MPC10E line card crashes without any known trigger. PR1627986
- On MX Series platforms with MPC10, and MPC11 line card, and MS-MPC/MS-MIC are used, if aggregated multiservices (AMS) interface is configured as next hop with equal cost multipath (ECMP), load balancing does not happen properly according to source IP hashing. PR1628076
- In a scaled subscriber service accounting scenario (~32K logical interfaces), if the `flat-file-profile` is configured with the `use-fc-ingress-stats` statement, the memory leak on pfed process might occur and if it crosses 80 percent of the total allocated memory of the process, it might crash. PR1628139
- On all Junos OS platforms, when unified ISSU is aborted, l2ald might not be able to read `issu abort` notification. The l2ald might be stuck in the `issu state` and will not process new events while in `issu state`. PR1629678
- The rpd process generates core file with the `warm-standby` configurations due to reference counting issues. PR1631871
- On MX platforms enabled with the `dynamic-profiles` for subscribers and the subscribers are configured over aggregate Ethernet interface with the `targeted-distribution`. When the child links of the aggregate Ethernet interface are removed and then added, it could lead to bbe-smgd crash in the backup Routing Engine. This in turn could affect the control plane subscriber services when the primary Routing Engine fails during such event. PR1633392

EVPN

- A few duplicate packets might be seen in an AA EVPN scenario when the remote provider edge device sends a packet with an IM label due to MAC not learned on the remote PE device, but learned on the AA local PE device. The nondesignated forwarder sends the IM-labeled encapsulated packet to the PE-CE interface after MAC lookup instead of dropping the packet, which causes the duplicate packets to be seen on the customer edge side. PR1245316
- The VXLAN OAM host bound packets are not throttled with DDoS policers. PR1435228
- On all Junos platforms in an EVPN-VXLAN to EVPN-MPLS stitching scenario, traffic loss might be seen with data forwarder (DF) changes when the traffic flows from VXLAN to MPLS. The traffic loss occurs till MAC IP ages out. PR1515096
- In a PBB-EVPN environment, the ARP suppression feature, which is not supported by the PBB might be enabled unexpectedly. This might cause the MAC addresses of the remote customer edge device not to be learned and hence the traffic loss might be seen. PR1529940
- VM core files are generated during evp-mpls script running while performing GRES. PR1580313

- EVPN-MPLS multihoming control MACs are missing after VLAN ID removal and adding it back to a trunk logical interface of one of the multihoming PE devices. This is not a recommended way to modify VLAN ID configuration. Always both multihoming PE devices needs to be in symmetric. PR1596698
- In an EVPN-VXLAN scenario in a datacenter and EVPN-MPLS in a WAN, and the stitching is done with an LT interface, then the bridge MAC table learning entries do not work as expected for EVPN-VXLAN routing instance. This might occur after the `restart interface-control` command is issued on gateways. PR1600310
- In EVPN VXLAN scenario, with the `proxy-macip-advertisement` statement is configured, a few ARP/ND/MAC entries might get missed. PR1609322
- MAC IP moves across L2-DCI is not updated in MAC-IP table of the gateway nodes. This problem happens only with the translation VNI when the MAC is moved from DC1 to DC2. VM moves across DC where there is no translate VNI configuration in the interconnect works as designed. PR1610432

Flow-based and Packet-based Processing

- Use 512 antireplay window size for IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). Hence there are no out-of-order packets with 512 antireplay window size. PR1470637

Forwarding and Sampling

- The `show firewall log detail` command shows the packet length for ICMPv6 as 0. PR1184624
- The configuration statement fast-lookup-filter with match condition is not supported in FLT hardware and might cause a traffic drop. PR1573350
- On MX Series platforms, when filter-based forwarding (FBF) is configured with the `next-interface` action and if the interface participating in the filter gets flapped due to any reasons, notification to update the filter action is initiated. It might result in FPC crash with core dump. This might be due to a rare timing issue. PR1622585

High Availability (HA) and Resiliency

- If you perform GRES with the interface em0 (or fxp0) disabled on the primary Routing Engine, then enabling the interface on the new backup Routing Engine might result in losing network access. PR1372087

Infrastructure

- If an interface is configured for single VLAN or multiple VLANs and all these VLANs of this interface have `igmp-snooping` enabled, then this interface will drop HSRPv2 packets. But if some VLANs do not have `igmp-snooping` enabled, then the interface works fine. PR1232403
- The following messages are seen during FTP: `ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock (No such file or directory).` PR1315605
- The `IFDE: Null uint32 set vector, ifd and IFFPC: 'IFD Ether uint32 set' (opcode 151)` error message is observed continuously in AD with base configurations. PR1485038
- The `show system processes detail` CLI command does not display CPU details under the CPU column. PR1588150

Interfaces and Chassis

- The CLI output for the `show interfaces transport pm otn current interface` command has a formatting issue with the interval range. The correct range information is returned in the commands XML message. The information can be displayed by redirecting the command output to `display xml`.PR1560533
- The issue is seen in a scaled setup with 296 LM sessions with iterator cycle time interval (100ms). It seems there is degradation in scale number (OAM packet rate at ~5500). At this qualified PPS, now LMR packet loss is observed but the functionality seems to be fine. To avoid LMR packet loss, reduce the scale number and keep the OAM packet value to less than 5500 pps. PR1561397
- Delay in application of CLI configuration by DCD when aggregated Ethernet interface members are configured via JET API. PR1621482
- On all Junos OS platforms, the **duplicate VLAN-ID on untagged interface gr-x/x/x.xxx: conflicts with unit 1** syslog message might be seen and dcd process might be crashed if the same VLAN is configured on the GRE tunnel interfaces. PR1633339

- In a Layer 3 VPN scenario with the `routing-options forwarding-table chained-composite-next-hop ingress l3vpn` statement is configured, if VRRP route tracking is used to track routes inside a VRF, and if such routes are with composite next hop, they might be marked as down even they are present in the VRF, hence the VRRP route tracking might not work properly. PR1635351

- On all Junos OS platforms, after upgrading, the VRRP state will not be correct and tracking routes of VRRP might show as unknown. The intended router might not be the VRRP master instead the peer router with less priority will be master. The route states are not correct because `route add` messages are not received at vrrpd after activation of the interface. When the interface is activated an interface route is created for the address configured on the interface, vrrpd receives the addition, then update the track route state accordingly. When this is not being received at the vrrpd, tracking routes might become unknown. PR1638378

Juniper Extension Toolkit (JET)

- The jsd process might take some time to detect abrupt termination of the socket at the collector or client side in certain cases. This can occur when flapping the interface on which the collector is connected to the router or when a firewall terminates the client port. In such cases, the client must wait for the connection termination to be detected, which could take around 1 hour, or restart the jsd process before being able to reconnect with the same client ID. PR1549044

Layer 2 Ethernet Services

- If the `request system zeroize` does not trigger zero-touch provisioning (ZTP), reinitiate the ZTP as a workaround. PR1529246

- On all Junos OS platforms configured as DHCP server or relay agent, the file system storage under /var directory might get filled up with DHCP event rate analyzer (ERA) logs, which is enabled by default and might result in other processes not having storage space to log details of router functionality. PR1617695

- On MX Series platforms, the jdhcpd process might crash and dump core files in a DHCP or DHCPv6 environment when the device is configured as a relay agent or server with the `active-leasequery` configuration. This might lead to subscriber termination and DHCP relay binding state of the terminating subscriber shows as `Release` state. PR1625011

MPLS

- On MPC7E line card, the BFD session flaps during unified ISSU and the issue is not seen frequently. PR1453705
- When we configure the `minimum-bandwidth`, the LSP is still resignalled with the previously configured minimum-bandwidth and not the currently configured minimum-bandwidth. PR1526004
- The single hop BFD sessions might flap sometimes after GRES in a highly scaled setup which have RSVP link or link-node-protection bypass enabled. This happens because the RSVP neighbor goes down sometimes after GRES if RSVP hellos are not received before neighbor time out happens. As a result of the RSVP neighbor goes down, RSVP installs a /32 route pointing to bypass tunnel which is required to signal backup LSPs. This route is removed when all LSPs stop using bypass after the link comes back. The presence of this /32 route causes BFD to flap. PR1541814
- The RSVP interface update threshold configuration syntax has changed between Junos OS Release 18.2X75-D435 and Junos OS Release 20.3X75-D10 to include curly braces around the threshold value. Upgrading and downgrading between these releases is not entirely automatic. The user must delete this stanza if configured before downgrading and then manually reconfigure. PR1554744
- If IS-IS-TE or OSPF-TE is enabled, but extended admin groups (which is configured under routing-options) are configured after the peer router advertises the extended admin groups, the LSP with extended admin groups constraints might fail to be established. PR1575060
- With the local reversion ON, there is a possibility of transit router not informing headend of RSVP disabled link when link is flapped more than once. As a workaround, remove the `local-reversion` configuration. PR1576979
- The `use-for-shortcut` statement is meant to be used only in SR-TE tunnels which use strict SPF (SSPF) Algo 1 prefix SIDs. If the `[set protocols isis traffic-engineering family inet-mpls shortcuts]` and the `[set protocols isis traffic-engineering tunnel-source-protocol spring-te]` are configured on a device, and if any SR-TE tunnel using Algo 0 prefix SIDs is configured with the `use-for-shortcut` statement, it might lead to routing loops or rpd process core files. PR1578994
- When a protected link goes down, MPLS gets tunnel local repair message from RSVP and trigger CSPF computation. Next, MPLS gets link protection information through RRO notification. If MPLS receives TED notification first before RRO notification, then CSPF computation fails. Because the link protection flag is not set, MPLS thinks it is an unprotected link and brings down the LSP. PR1598207

Network Management and Monitoring

- When the ephemeral instance is deleted, physical files related to the instance is not deleted and the content of the file will remain as it is and might cause the device to behave uncertain. PR1553469
- The shm-rtsdbd daemon generates core files when the services configuration is deactivated or activated. PR1610594

Platform and Infrastructure

- On MX Series platforms with MPC7, MPC8, MPC9 line card or MX-204 and MX-10003, when the packets which exceed the MTU and whose DF-bit is set go into a tunnel (such as GRE and LT), they might be dropped in the tunnel egress queue. PR1386350
- The following error message is observed during unified ISSU: `Async XTXN Error PPE/Context 9/13 @ PC 0x6f77: sampling_li_launch_nh`. The traps are the result of PPE commands injected from the host. One possible reason might be Layer 2 BD code, which is trying to decrement BD MAC count in the data plane. It is unlikely that there is a packet loss during this condition. This could happen during unified ISSU and this might be due to a problem with the ISSU counter morphing used for LU-based cards, where certain counters are not disabled or disabled too late during unified ISSU. PR1426438
- Arrival rates are not seen at system level when the `global-disable fpc` is configured. PR1438367
- Loss of traffic on switchover when using filter applied on logical interface. PR1487937
- When GRES and NSR functionality with VXLAN feature, the convergence time might be slightly higher than expected for Layer 2 domain to Layer 3 VXLAN. PR1520626
- When the DHCP relay mode is configured as `no-snoop`, we observe the offer gets dropped due to incorrect ASIC programming. This issue happens only while running DHCP relay on EVPN-VXLAN environment. PR1530160
- On MX Series platforms with XM chipset based line card installed, when the line card experiences the `XMCHIP_CMERROR_DDRIF_PROTECT_WR_RD_SRAM_RUNN_CHKSUM` CM error, the disable-pfe action will be involved. This issue causes the Packet Forwarding Engine to be disabled and traffic lost. PR1568072
- On MX Series platforms, FPC gets restarted and thereby disrupting traffic when there is an out-of-order filter state. This issue might be seen only in back-to-back GRES in more than 40 to 50 iterations. PR1579182
- Ethernet-output-bytes are not in expected range while verifying Ethernet MAC level with both IPv4 and IPv6 traffic for VLAN tagged interfaces. The issue is due to output byte count not getting

updated properly. The script log shows that there is no packet loss and there is no functional impact. PR1579797

- MS-PIC RPM probes with large data-size is failing at random. PR1602508
- On all MX Series platforms, in a rare case when CoS classifier binding message received before logical interface family creation message to Packet Forwarding Engine, traffic might be classified with default classifier instead of custom classifier. Due to this, traffic might not be classified and mapped to the right queue resulting in inappropriate CoS treatment for the traffic. PR1619630
- On MX Series platforms, during reboot, the aggregated Ethernet logical interfaces are first added, then deleted and again added, this flapping causes corner case where the filter attachment ipc has older aggregated Ethernet logical interface index on which the filter bind fails. Filter will not be attached to the interface, so any filter related service will not work. PR1614480

Routing Policy and Firewall Filters

- On all Junos OS platforms with the `set policy-options rtf-prefix-list` configured, if you upgrade to a specific version, the device might fail to validate its configuration, which eventually causes rpd to crash unexpectedly due to a software fault. PR1538172

Routing Protocols

- While interoperating with other vendors in a draft-rosen multicast VPN, by default Junos OS attaches a route target to multicast distribution tree (MDT) subsequent address family identifier (SAFI) network layer reachability information (NLRI) route advertisements. But some vendors do not support attaching route targets to the MDT-SAFI route advertisements. In this case, the MDT-SAFI route advertisement without route-target extended communities are prevented from propagating if the BGP route-target filtering is enabled on the device. PR993870
- If delegated BFD sessions flap continuously, packet buffer memory might be leaked. The automatic memory leak detection process reports this within the syslog once a certain threshold is reached. The following error is displayed: `fpc7 SHEAF: possible leak, ID 8 (packet(clones)) (10242/128/1024) on MX-MPC or fpc4 SHEAF: possible leak, ID 9 (packet(clones)) (255/1/5)`. Note that BFD sessions operating in centralized mode are not exposed. PR1003991
- Certain BGP traceoption flags (for example, open, update, and keepalive) might result in (trace) logging of debugging messages that do not fall within the specified traceoption category. This results in some unwanted BGP debug messages being logged to the BGP traceoption file. PR1252294

- LDP OSPF are in synchronization state because the IGP interface is down with the `ldp-synchronization` enabled for OSPF. As per the current analysis, the IGP interface goes down because although LDP notified OSPF that LDP synchronization was achieved. OSPF is not able to take note of the LDP synchronization notification because the OSPF neighbor is not up yet. PR1256434

- In rare cases, RIP replication might fail as a result of performing NSR Routing Engine switchovers when the system is not NSR ready. PR1310149

- On MX Series platforms, the following unexpected log message will appear if the `show version detail` or `request support information` CLI command is executed: `user@host> show version detail *** messages *** Oct 12 12:11:48.406 re0 mcsnoopd: INFO: krt mode is 1 Oct 12 12:11:48.406 re0 mcsnoopd: JUNOS SYNC private vectors set`. PR1315429

- SCP command with routing Instance -JU is not supported. PR1364825

- TILFA backup path fails to install in LAN scenario and also breaks SR-MPLS TILFA for LAN with more than four end-x SIDs configured per interface. PR1512174

- Conformance issues with draft-ietf-idr-bgp-ext-opt-param. In previous versions of RFC 9072 (that is, draft-ietf-idr-bgp-ext-opt-param), the required optional-parameter length is 255 in order to trigger the updated behavior. Later editions of the internet draft permitted non-zero optional parameter length values to be used. PR1554639

- Due to behavior change, if there is no IFA present in the interface, we do not encode the router ID in the hello packet by default. In current scenario between R1 and R2, we do not have any inet or inet6 address set for interfaces forming the adjacency in question. Then, in the `show isis adjacency detail` command output, we do not see IPv4 or IPv6 address and it is shown that the adjacency is missing an IP address. PR1559079

- With maximum number of logical interfaces (4000 GRE tunnels per Packet Forwarding Engine) with the following configuration:

  1. family inet, associated source, and destination for each tunnel.

  2. Configure the `allow-fragmentation` statement on one endpoint of the tunnel and configure the `reassemble-packets` on the other endpoint of the tunnel.

  With the above configuration, if you do `deactivate chassis fpc slot`, SLIP messages are observed. PR1581042

- On all Junos OS platforms running BGP with Layer 2 VPN, kernel crash might be observed. PR1600599

- When MPLS traffic engineering and the rib inet.3 protect core statement is enabled, then the transport routes in inet.3 will not be used for route resolution. PR1605247

- On all Junos OS platforms, traffic drops when incorrect VPN labels are allotted. When there is a change in the next hop by BGP policy, the traffic is still forwarded to the old label. This leads to traffic drops for prefixes sending traffic to the old next hop. PR1617691
- On all Junos OS platforms, if an aggregate route is configured under routing options, and if the `from aggregate-contributor` is used for many contributing routes (for example, more than 250-300 routes), the policy for these contributing routes might not work properly when the policy is exported. Due to this issue, the contributing routes might not be advertised properly. PR1629437
- On all Junos OS platforms with multicast setup, the multicast forwarding cache might not get updated after deactivating the `scope-policy` configuration. This might result in the PIM register process to be incomplete and further multicast traffic to be dropped. PR1630144
- When IS-IS database is cleaned, rpd crash might be observed. PR1631738
- On all Junos OS platforms that support NSR, when the `switchover-on-routing-crash` is enabled, the rpd process crash will lead to Routing Engine switchover. In a highly scaled environment (about 15~19 million BGP routes), BGP session which is still sending update packets of size more than 2000 might flap even when NSR is enabled. This might lead to loss of traffic till the BGP session converges after the flap. This does not happen always but happens sporadically. The switchover can be either due to the rpd process crash or when switchover is performed manually. PR1632132
- On all platforms with IS-IS multiple areas scenario, if the `flood-group` statement is enabled, IS-IS databases might not get synchronized between areas after clearing the IS-IS database or making the database change in any other way. This is because when the LSP is fragmented, only the first packet has the area ID list (for flood-group matching), while the rest of the fragmented LSPs do not have that list, which will result in these packets not being flooded, so that IS-IS will not work properly. PR1633858
- On all Junos OS platforms running BGP, when a specific route is received from multiple places under a VRF, multipath route is getting formed even though the BGP route selection algorithm has the active route with higher local preference. Once multipath is formed, the traffic forwarding happens based on that, and it might result in some traffic going to an unwanted path. PR1635009
- In a scenario where the single hop BFD of BGP, when multiple addresses of the same subnet are configured on the interface of the BFD session, the BFD session might go down. PR1635700

Services Applications

- In an L2TP environment on L2TP LAC (L2TP access concentrator), a few L2TP tunnels might be stuck in downstate and might not be able to reestablish if the bbe-smgd process is restarted when these tunnels go down, which might impact the end customer to lose connectivity. PR1629104

- On all MX platforms that support enhanced subscriber management (next generation subscriber management) with L2TP subscriber scenario, L2TP subscribers might get stuck in terminating state if the L2TP subscribers try to login. PR1630150

Subscriber Access Management

- When performing unified ISSU from earlier releases to certain releases on MX Series platforms, accounting messages for new service on existing subscribers will have corrupted class attribute value, which might be rejected by RADIUS server. As a result, new service on existing subscribers might not get created. PR1624066
- When the extensible subscriber services manager (ESSM) service is getting created on existing subscriber session, the class attribute is incorrectly formed. This happens when RADIUS sends class attribute in access accept messages after performing unified ISSU. PR1626718

User Interface and Configuration

- The mgd process generates core file upon simultaneous rollback command in two different terminals of same router. It is a rare and corner case and is a timing issue. If this happens, the CLI session ends abruptly. PR1554696
- In an EVPN-VXLAN scenario, mgd process generates core file when executing image upgrade command. The issue is seen on Virtual Chassis only, which can be avoided with a simple workaround by providing a valid package during upgrade command. PR1557628
- After several ephemeral commits, interface configurations might get stuck and might not get updated on all Junos OS platforms. PR1598123
- When performing commit check for the firewall and interface related configurations, if an operator uses the Ctrl+C to abort it, the dfwc and dcd might crash after performing another commit check. This issue will happen only with those daemons that follow the message-based commit check model (such as dfwc, dcd, rdmd, and fwa), and has no impact on other daemons. PR1600435

VPNs

- In some scenario (for example, configuring firewall filter), routers might show obsolete IPsec SA and NHTB entry even when the peer tear down the tunnel. PR1432925

- Incorrect st0 logical interface deletion at spoke when multiple VPNs negotiate same destination address as TS. The general trigger is that when multiple VPNs configured have the traffic selectors which have the same remote-ip or subnet. And if one of the tunnels go down, the incorrect st0 route gets deleted. PR1601047

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for MX Series routers.

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

Resolved Issues: 21.1R3

IN THIS SECTION

Application Layer Gateways (ALGs)

- ALG traffic might be dropped. PR1598017
- On MX Series routers, the flowd process crashes if the SIP ALG is enabled and specific SIP messages are processed. PR1604123
- JFlowd core files are observed if the SIP ALG is enabled and a specific SIP packet is received. PR1615438

Class of Service (CoS)

- Transit packets from local to remote VTEP might get punted to CPU and cause DDoS events. PR1489233
- Child mgd processes might get stuck when multiple sessions continuously ask for interface information. PR1599024
- Traffic loss might be observed if the `per-unit-scheduler` is configured on aggregated Ethernet interface. PR1599857

EVPN

- Traffic loss might be seen under an EVPN scenario when MAC-IP moves from one CE interface to another. PR1591264

- Transit traffic gets dropped post disabling one of the PE-CE links on a remote multihome PE in an EVPN-MPLS AA setup with dynamic-list next hop configured. PR1594326
- EVPN might not work properly in multihoming setup. PR1596723
- Adjusting mac-ip binding table might only be activated by rebooting l2ald. PR1599305
- Bridge mac-table learning entries might not be as expected for the EVPN-MPLS routing instance. PR1600310
- The device announces router-mac, target, and EVPN VXLAN community to BGP IPv4 NLRI. PR1600653
- Missing MAC address entries in the EVPN mac-table despite the presence of the corresponding Type 2 route. PR1611618

Forwarding and Sampling

- Logical interface statistics for aggregated sonet displays double value than the expected. PR1521223
- The IPv6 filter for family bridge cannot be referenced onto logical interface. PR1598530
- The snmpwalk might not get polling the MIB for dual stack interface. PR1601761
- More than 5 minutes delay in getting the response for the `clear interfaces statistics all` command with RIB scale configuration. PR1605544
- Commit is allowed even if firewall filter is not applied to the FPC. PR1618231

General Routing

- On MX10003 routers, despite of having all AC low or high PEM, the `Mix of AC PEMs` alarm is raised. PR1315577
- Traffic loss is observed after FPC reboot in a scaled scenario when reverting RLT during FPC resync. PR1394026
- Junos OS does not provide any logging for non-SNI sessions.PR1442391
- Inaccurate allocated memory for `nh` and `dfw_rulemask` under kernel might be observed. PR1475478
- Next hops are not programmed correctly after Virtual Chassis global switchover PR1518467
- The PKI CMPv2 client certificate enrollment does not work when using root-CA. PR1549954
- The Virtual Chassis port might not come up after upgrading. PR1555741

- MPLS Jflow packets are dropped on the MPLS interfaces. PR1559390
- The untagged traffic routed over `native-vlan` might be dropped. PR1560038
- The MX150 router might reboot after performing the `request system snapshot recovery` command. PR1565138
- Local privilege escalation and denial of service. PR1568654
- Wi-Fi mPIM is reaching out to NTP and DNS servers. PR1569680
- NG logging profiles do not function properly with unified-policies. PR1570105
- The PDB pull or synchronization might fail during unified ISSU. PR1570841
- BFD sessions over VTEP might fail. PR1571417
- Switchover to backup Routing Engine if IFF rpd is NSR ready and then crashed. PR1571914
- High CPU usage might occur on rpd for routes that use static subscriber. PR1572130
- DCPFE or FPC crash might be observed if the ARP MAC move happens. PR1572876
- DS-Lite throughput degradation might be seen on MS-MPC. PR1574321
- The chassisd process might crash on all Junos platforms that support Virtual Chassis or Junos fusion. PR1574669
- The `CHASSISD_FRU_IPC_WRITE_ERROR: fru_send_msg: FRU GNF 2, errno 40, Message too long` error might appear periodically in the chassisd logs. PR1576173
- The OSPF session over IRB might not come up in an EVPN-VXLAN scenario. PR1577183
- Kernel crash might be observed on the backup Routing Engine after GRES. PR1577799
- The MPC7E, MPC8E, MPC9E, and MPC11E line cards might be stuck in `Unresponsive` state in a Junos node slicing setup. PR1580168
- The bbe-smgd process crashes if an unsupported configuration exists and a PPPoE client sends a specific message. PR1580528
- VM core might be seen after adding and deleting the logical interface of the static interface in the next generation subscriber management subscriber scenario. PR1581260
- The pciephy and firmware download is not working after a migration to 6.5.19. PR1582244
- A vulnerability in the Juniper agile license client might allow an attacker to perform remote code execution (RCE). PR1582419

- Traffic drop might be observed on MX Series platforms with SPC3 in the DS-Lite scenario. PR1582447
- Load balancing is not working correctly on AMS interfaces for CGNAT traffic on MX USF mode with SPC3. PR1582764
- The bcmd process might crash on the MX150 router. PR1583281
- The firewall filter does not get programmed after you delete a large filter and add a new one in a single commit. PR1583440
- Layer-2 multicast VXLAN instance is down as local VTEP logical interface is not associated to EVPN instance. PR1584109
- Secure web proxy continue sending DNS query for unresolved DNS entry even after the entry is removed. PR1585542
- MX104 routers might become unresponsive if the out-of-band management port receives a flood of traffic. PR1585829
- A high rate of small packets could cause CPU hogging and the firmware crash in MPC5E and MPC6E line cards. PR1587551
- The MVPN traffic loss might occur due to missing of the flooded multicast next hop.PR1587054
- The na-grpc process might crash and existing telemetry connections will be disconnected. PR1587956
- The dcpfe might crash when loading EVPN with VXLAN configuration on the setup. PR1588637
- The MPLS traffic might not be forwarded after the aggregate interface flaps. PR1589840
- The NAT service might not happen after performing AMS switchover or deactivating and activating NAT service. PR1590890
- Some logical interfaces might go down under logical tunnel due to the limited number of MAC addresses in a pool. PR1591853
- AMS warm standby with deterministic NAT functionality might not work properly. PR1592437
- The l2cpd-agent might go unresponsive after starting telemetry service. PR1592473
- Using the BITS interface from backup Routing Engine for clock recovery might not work. PR1592657
- The TCP connections to the telemetry server might be stuck in `CLOSE_WAIT` status. PR1593113
- The IPv6 neighbor might remain unreachable in a VRRP for an IPv6 scenario. PR1593539
- J-web deny log nested-application is UNKNOWN instead of specific application. PR1593560

- The dcpfe process might crash in an EVPN-VXLAN scenario. PR1593950
- Executing `set interfaces disable` command on an interface in the set [ge-0/0/8~11] might cause traffic to stop on another port in this set. PR1593983
- In an EVPN-VXLAN scenario, the label field for Type-1 route is not required but it is assigned 1 instead of 0. PR1594981
- Memory usage continuously increases on backup chassis if the subscriber service is enabled. PR1595238
- The interface down might be delayed after performing the `set interface interface name disable` command. PR1595682
- The Packet Forwarding Engine wedge might be seen if received many IPv4 packets that need to be fragmented. PR1596100
- The DCI InterVNI and IntraVNI traffic might get silently dropped and discarded in gateway node due to the tagged underlay interfaces. PR1596462
- The mcscnoopd might crash during deleting and adding Layer 2 forwarding configuration after performing unified ISSU. PR1596483
- The l2alm fails to send IPC message to the l2ald which might cause the FPC to crash. PR1596615
- Traffic loss might happen periodically in MACsec used setup if Routing Engine is working under a pressure situation. PR1596755
- The SR-TE tunnel initiated from a non-Juniper PCE might fail. PR1596821
- The bbesmgd process generates core file after Routing Engine goes down. PR1596848
- The MAC/IP withdraw route might be suppressed by the rpd in an EVPN-VXLAN scenario. PR1597391
- The mspmand process might crash if memory leak issue occurs on MX Series platforms with MS-MPC and MS-MIC. PR1597624
- Deletion of MACsec configuration on a logical interface is not taking effect. PR1597848
- Subscriber management daemons might continuously generate core file and shutdown with Routing Engine sensors invalid configuration. PR1598351
- The packet loop might be seen after receiving the PCP request packets which are destined to softwire concentrator address. PR1598720
- Subscriber might not be able to come up on an aggregated Ethernet interface. PR1598726
- Component sensor does not export logs. PR1598816

- The rpd process generates core file if BGP update tracing is configured and an update containing a malformed BGP SR-TE policy tunnel attribute is received. PR1598850
- The l2ald process might crash due to memory leak when all active interfaces in a VLAN are unstable. PR1599094
- On MX SPC3 services card, ICMP protocol is not detected and does not allow user to modify inactivity-timeout values. PR1599603
- The multiservices card does not drop the TCP ACK packet received as a reply to the self-generated TCP keepalive. PR1600619
- Interface flap might trigger major alarm causing disble-pfe action when high priority scheduler is configured. PR1601049
- Duplicate address detection (DAD) flags can be seen for IRB interfaces after removing configuration and restoring, which might lead to blocking the traffic. PR1601065
- The BBE-SMGD core files are found at `bbe_dequeue_and_deliver bbe_process_work_queues bbe_smd_main_post_dispatch`. PR1601203
- Unable to commit configuration due to error check-out failed for mobility process. PR1601785
- Traffic might be dropped at NAT gateway if EIM is enabled. PR1601890
- The IPv6 traffic might be impacted when an IPv6 route resolves over a dynamic tunnel. PR1602007
- Some EVPN MAC IP entries might get stuck and does not age out. PR1602010
- A few line cards might not come up online with increased bandwidth mode. PR1602080
- Under certain scaling scenarios in an EVPN-VXLAN scenario, l2ald might abort and recover. PR1602244
- After upgrading, configured firewall filters might be applied on incorrect interfaces. PR1602292
- Specific packets over VXLAN cause FPC memory leak and ultimately reset. PR1602407
- J-Flow-syslog for CGNAT might use 0x0000 in IPv4 identification field for all fragments. PR1602528
- An l2cpd memory leak can occur when specific LLDP packets are received leading to a DoS. PR1602588
- CRL failing to download causes a memory leak and ultimately a DoS. PR1602815
- Traffic might be dropped in a Virtual Chassis scenario when the firewall filter is configured. PR1602914

- The Packet Forwarding Engine might be disabled by a detected major CMERROR event while ungracefully removing the MIC from MPC2E-3D-NG or MPC3E--3D-NG. PR1602939
- When using J-Web with HTTP, an attacker might retrieve encryption keys via person-in-the-middle attacks. PR1603199
- Packet loss might be seen on the filter-based GRE deployments. PR1603453
- The PTP slave might stay in holdover and does not process the PTP packets. PR1603483
- Traffic loss might be seen on the device due to the continuous errors happening on fabric healing process (FHP). PR1603499
- The fxpc core files are generated when the NSSU performed with MACsec configuration. PR1603602
- VRRP and BFD might flap on IRB interface on MPC10 and MPC11 line cards. PR1604150
- NPC logs are seen when vrf localisation is enabled. PR1604304
- Interface hold-time up does not work on MX150 routers. PR1604554
- Authd process might crash and generates core files when trying to connect to Juniper Secure Connect. PR1604616
- GRE tunnel might flap when `hierarchical-scheduling` is configured. PR1605189
- The interface on MCP3-NG HQoS and MPC7E flaps continuously after enabling LACP on aggregated Ethernet interface. PR1605446
- The MPLS transit router might push an extra entropy label to the LSP. PR1605865
- Multicast streams might stop flooding in a VXLAN setup. PR1606256
- In a scenario with the `dhcp-security` and `option-82` are configured, the jdhcpd crashes upon receipt of a malformed DHCP packet. PR1606794
- After an FPC oversubscription, new subscribers might not be able to connect. PR1607056
- The TCP traffic might be dropped on source port range 512 to 767 when the FlowSpec IPv6 filter is configured. PR1607185
- In a subscriber management scenario, under a rare condition, the Routing Engine reboots and generates a vmcore file. PR1607282
- Commit related to dynamic profile configuration changes might fail upon executing the `request vmhost reboot routing-engine both` command. PR1607494

- On MX104 routers, if the SFP-T optic connected interface negotiates a speed other than 1G and is part of an aggregate interface, the interface's negotiated speed will not be shown after the interface-control (dcd) daemon restart or a Routing Engine switchover. PR1607734
- Memory leaks might be observed on the l2cpd process when performing certain LLDP operations. PR1608699
- The GNMI set RPC does not work with multiple operations. PR1609436
- DHCP subscribers over PWHT might be dropped upon GRES after the system reboot. PR1609818
- The single VLAN tagged subscribers might fail to reconnect through dynamic VLAN over PS interface. PR1609844
- The authd process and RADIUS might have stale L2BSA subscriber entries. PR1610476
- Traffic loss might be observed if dot1X is configured with 'supplicant multiple' and authenticated user from radius is in single supplicant mode PR1610746
- MACsec session might be dropped due to one way congestion. PR1611091
- Erratic behavior might be seen on platforms using MPC line cards after unified ISSU is performed. PR1611165
- Inter VLAN connectivity might be lost in an EVPN-VXLAN with CRB topology. PR1611488
- The CPU of the routing protocol engine gets stuck at 100 percent. PR1612387
- The B4 client traffic will be dropped on MX-SPC3 based AFTR in DS-Lite with EIM activated CGNAT scenario. PR1612555
- The l2ald process generates core file during routing instance configuration change. PR1612738
- Memory might be exhausted when both the BGP RIB sharding and the BGP ORR features are enabled. PR1613104
- Traffic loss might be observed due to the shaping rate be adjusted incorrectly in a subscriber environment. PR1613126
- The `enhanced-hash-key` might not take effect when configured with the `forwarding-options`. PR1613142
- The IGP routing updates might be delayed to program in Packet Forwarding Engine after interface flaps in a scaled BGP routes environment. PR1613160
- The rpd process might crash in a BGP RIB sharding scenario. PR1613723
- Any irrelevant configuration changes might trigger NAT routes flap on MX routers in USF mode. PR1614688

- Line cards might be unstable due to the continuous growing memory usage. PR1614952
- The l2ald process might crash in an EVPN scenario. PR1615269
- Traffic drop might occur when huge number of EIM mappings are created and deleted continuously. PR1615332
- Slow memory leak (32 bytes each time) of the rpd might be seen. PR1616065
- The `show subscribers accounting-statistics` and the `show services l2tp session interface asi0.xx statistics` might not work on LNS with asi- interfaces. PR1616454
- The dual Routing Engine system might not be GRES ready after backup Routing Engine reboots in a subscriber management environment. PR1616611
- Memory leak might be seen when LLDP is configured. PR1617151
- On MPC8E line card in 1.6T bandwidth mode might not work correctly. PR1617469
- Traceroute packets might get dropped in SFW service set when other service sets with asymmetric traffic processing are also enabled on the same MS-MIC and MS-MPC. PR1617830
- GMC clock class is seen transmitted for an additional 16 seconds after the PTP source switches from one line card to another. PR1618344
- The CGNAT traffic loss might be seen after cleaning the large scaled CGNAT sessions in an MS-SPC3 based inter-chassis high availability scenario. PR1618360
- The clksyncd process might crash and PTP/SyncE might not work. PR1618929
- Incorrect NHG information reported in telemetry data for LSP after addition or deletion of low priority tunnels with the regular LSPs. PR1619011
- Traffic might be dropped when the RSVP is configured with the `mtu-signaling`. PR1619510
- The bbe subscriber access services might be stuck during rebooting the one redundancy line card of RLT interface. PR1620227
- On MX480 routers, output packet drop is observed while verifying services PCEF subscribers. PR1620421
- OAM CFM session does not come Up if ERPS configured and CFM control traffic uses the same VLAN as ERPS control traffic. PR1620536
- High wired memory utilization might be observed if GRES is enabled. PR1620599
- The EVPN type 5 routes might not be installed. PR1620808

- Static subscribers session might get stuck in initializing state after ungraceful routing engine switchover. PR1620827
- All ports from the same Packet Forwarding Engine go down at the same time causes mqchip_disable_ostream timeout and triggers host loopback path wedge. PR1621286
- Traffic loss might be seen on the new master Routing Engine post GRES. PR1621696
- The aggregated Ethernet member link might not be correctly populated on the Packet Forwarding Engine after FPC restart on MX Series platforms. PR1624772
- Invocation of NETCONF get command will fail if there are no Layer 2 interfaces in the system. PR1622496
- The chassisd memory leak might be seen after adding or removing an interface configuration. PR1623273
- Introduce a new `show task scheduler-slip-history` command to display number of scheduler slips and last 64 slip details. PR1626148
- The chassisd process might crash on MX104 routers. PR1626486
- Specific packets over VXLAN cause FPC reset. PR1625292
- The gNMI set RPC might fail when multiple values within a single gNMI SetRequest is used for telemetry interface. PR1625806
- The bbe-smgd might crash on backup Routing Engine after unified ISSU or GRES. PR1626091
- Some interfaces might not come online after line card reboot. PR1626130
- Layer 3 traffic failures are observed continuously with PDT mclag configuration. PR1627846
- The EAPoL packets over l2circuit might get dropped at the tunnel start. PR1628196
- The kmd process might crash and generates core file every few minutes on MX Series platforms. PR1630070
- The LLDP packets might be sent with incorrect source MAC for RETH/LAG child members. PR1630886
- The kmd process might crash because the pkid requested memory leak happens on MX Series platforms. PR1631443
- When you add and remove the VLANs, the traffic loss might be observed. PR1632444
- Data might not be exchanged via EVPN-VXLAN domain. PR1635347

High Availability (HA) and Resiliency

- Memory leaking might occur on backup Routing Engine when ksyncd is in inconsistent state and had encountered an initialization error. PR1601960
- When MTU is configured on an interface, a rare ifstate timing issue might occur at a later point, resulting in ksyncd process crash on backup Routing Engine. PR1606779

Infrastructure

- The fxpc process might crash and generate core files. PR1611480

Interfaces and Chassis

- Traffic might be interrupted while adding xe- and ge- interfaces as member of aggregated Ethernet interface bundle. PR1569399
- The ARP resolution failure might occur during VRRP failover. PR1578126
- The dcd process might crash after performing Routing Engine switchover, reboot, or management interface configuration change. PR1587552
- The dcd process might crash after removing aggregated Ethernet logical interface from the targeted distribution database. PR1591032
- Duplicate source and destination pair check is done only across the same tunnel encapsulation type for FTI. PR1599266
- The dcd process might crash and FPC might be stuck in ready state on MX Series platforms. PR1601566
- The aggregated Ethernet interface might flap upon configuration changes. PR1602656
- Memory leak on dcd process occurs when committing configuration changes on any interfaces in a setup with AMS interface configured. PR1608281
- On MX960 routers, the following syslog messages are found: `dcd[40867]: %DAEMON-5: lo0 family maximum labels is non-adjustable`. PR1611098
- Commit check failure might happen if similar interfaces are configured under VRRP group. PR1617020
- The subscribers might be deleted when the `host-prefix-only` statement is configured on the underlying-interface in GRES scenario. PR1630229

J-Web

- J-Web allows a locally authenticated attacker to escalate their privileges to root. PR1594516

Junos XML API and Scripting

- Certificate validation is skipped when fetching the system scripts from a HTTPS URL. PR1542229

Layer 2 Ethernet Services

- Making configuration changes with the `apply-group` add or delete associated with DHCP might result in client connection failure. PR1550628
- The subscriber login might fail on backup BNG running ALQ and redundancy services will not be available. PR1583445
- The jdhcpd process might be stuck at 100 percent CPU usage. PR1585493
- The DHCP client might be offline for about 120 seconds after sending the DHCPINFORM message. PR1587982
- The delegated prefix IPv6 address is missing in the accounting stop messages. PR1588813
- The DHCP ALQ queue might get stuck and cause subscriber flap. PR1590421
- The jdhcpd process crashes upon receipt of a specific DHCPv6 packet. PR1594371
- The jdhcpd process might crash under certain conditions. PR1603992
- The DHCP leasequery fails to restore binding when the reply is received over IRB interface. PR1611111
- The jdhcpd process crashes upon receiving a specific DHCP packet. PR1618977
- The rpd scheduler might continuously slip after GRES when there are 7000 DHCP clients in a subscriber management environment. PR1625617
- The non-DHCPv4 BOOTP protocol packets might not be processed if enhanced subscriber management is enabled. PR1629172

MPLS

- The D-CSPF node segment label is unresolved when node index 0 is configured. PR1564169
- The rpd process generated core file in backup Routing Engine at `mirror_process_recvd_data_queue` with mLDP NSR configuration. PR1594405

- The LDP replication session might not get synchronized when `dual-transport` is enabled. PR1598174
- The rpd process might crash with LSP external controller configuration. PR1601763
- The VPLS connection might get down if the `dual-transport` statement is configured. PR1601854
- The RSVP detour LSP might fail to come up when an LSR in the detour path goes down. PR1603613
- LDP P2MP traffic might be interrupted post GRES. PR1609559
- The rpd process might crash on standby_re LDP module when the VPLS mac-flush is enabled on peer by default or configuration. PR1610638
- The rpd core files might generate for a few value configurations of signaling bandwidth on container LSP. PR1614248
- Protected LSP goes down when strict hops and link protection are configured. PR1616841

Multicast

- Intermittent p2mp traffic drop might be seen in an MVPN scenario. PR1608311

Network Address Translation (NAT)

- Services NAT mappings and sessions are incorrect while checking the SIP sessions from public to private and RTP from private to public. PR1577922
- The master-eventd process might go down when syslog configuration is misconfigured. PR1611885
- Syslog messages might be lost partially in case of lots of messages generated to eventd. PR1612535
- After receiving a specific number of crafted packets, snmpd segmentation fault (SIGSEGV) requires a manual restart. PR1613874

Network Management and Monitoring

- SNMP reflects outdated ARP entries. PR1606600

Platform and Infrastructure

- The ppmd process might crash after an upgrade. PR1335526
- The SPC3 might not come up after the system reboot. PR1555904
- Configuration changes of chassis FPC might cause traffic loss for the MPC7, MPC8, MPC9 or similar line cards. PR1585576

- The system generates an audit core file while changing TACACS and login user passwords. PR1589953
- The subscribers might not come online after interface flaps on MX Series platforms. PR1591905
- Upon receipt of specific sequences of genuine packets destined to the device, the kernel will crash and restart. PR1595649
- The **XMCHIP_CMERROR_PT_INT_REG_PCT_PAR_ERR (0x70296)** error might be observed on MPC5 line card which triggers Packet Forwarding Engine disable. PR1597953
- The VLAN tagged traffic might be dropped with service provider style configuration. PR1598251
- The service filter might get incorrectly programmed in the Packet Forwarding Engine due to a rare timing issue in an enhanced subscriber management environment. PR1598830
- There might be FPC core file and packet drop in an EVPN-VXLAN scenario. PR1600030
- Traffic through one SPU might stop with potential packet drop issue with alarm as FPC Major Errors raised due to the **PIC_CMERROR_TALUS_PKT_LOSS** error. PR1600216
- The kernel core file might be seen if you restart BGP connections after deleting the BGP authentication. PR1601492
- The ZTP service might not work and the image installation fails. PR1603227
- The FPC might crash if the `flow-table-size` is configured on MX Series platforms. PR1606731
- Multicast traffic is dropped when it is forwarded over VPLS via IRB. PR1607311
- FPC might crash due to MAC move between two interfaces under same bridge domain. PR1607767
- Degraded traffic processing performance might be observed in case of processing very high PPS rate traffic. PR1619111
- Accounting and auditd process might not work on a secondary node. PR1620564
- MX Series with MPCs and MICs might crash when Packet Forward Engine memory is hot-banking. PR1626041
- Unrealistic service accounting statistics might be reported due to the firewall counter corruption. PR1627908

Routing Policy and Firewall Filters

- The interface-routes rib-group policy does not work as expected in the VXLAN scenario. PR1537306
- The configuration check might fail if more than 8 FCs are configured and CBF is enabled. PR1600544

- Evaluation of inet-vpn route filters might not work with /32 exact statements for BGP flowspec routes. PR1618726

Routing Protocols

- EVPN database might not be populating when moving EVPN instances from the master to logical systems. PR1504590
- Short multicast packets drop using PIM when multicast traffic received at a non-RPT or non-SPT interface. PR1579452
- The rpd process might crash in a BGP multipath scenario if the single hop EBGP peer goes down. PR1585265
- The rpd process might crash when BGP RPKI session record-lifetime is configured less than the hold-time. PR1585321
- Traffic drop might occur on link flap when IS-IS is configured. PR1585471
- The rpd process might crash post GRES. PR1590912
- The rpd process might crash if the BGP peer flaps. PR1592123
- The IPv4 static route might still forward traffic unexpectedly even when the static route configuration has already been deleted. PR1599084
- Some BFD sessions might stuck in Init state after the FPC restart. PR1599431
- The OSPFv3 session might go into INIT state upon receipt of multiple crafted packets from a trusted neighbor device. PR1599491
- The rpd core might be observed due to memory corruption. PR1599751
- Some routes might get incorrectly programmed in the forwarding table in the kernel with next hop installed as DEAD. PR1601163
- The rpd process might get stuck at 100 percent in an OSPFv3 scenario. PR1601187
- Packet drop might be seen when changing INET MTU for MPLS enabled interface in a IS-IS SPRING scenario. PR1605376
- On MPC10E line cards, the rpd process generates core files after deactivating and activating interfaces: `rt_table_flash_job_cancel,rt_instance_set_lsi_ifl_data_shard,rt_flash_all_internal`. PR1605620
- The BGP replication might be stuck in the `InProgress` state. PR1606420
- Multicast traffic might be duplicated on subscriber interface on MX Series platforms. PR1607493

- The rpd process might crash with the telemetry setup. PR1607667
- With the `rib-sharding` enabled, any commit will flap all the BGP sessions with 4 byte peer-as (AS number 65536 or greater). PR1607777
- The rpd process might crash after a commit if there are more than one address in the same address ranges configured under the `bgp allow`. PR1611070
- The interface might receive multicast traffic from a multicast group which it is not interested in. PR1612279
- The rpd process might crash on all Junos platforms. PR1613384
- Undesired protection path might get selected for some destination prefixes. PR1614683
- The memory leak on the rpd process might be observed after running the `show route` CLI command. PR1615162
- BFD sessions flapping might occur after performing GRES. PR1615503
- The incorrect BGP path migth get selected even when a better or preferred route is available. PR1616595
- The rpd process might crash and restart when NSR is enabled. PR1620463
- Time delay to export prefixes to BGP neighbors might occur post applying peer-specific BGP export policies. PR1626367
- Multipath route with list next hop which has indirect next hop as members fails into BGP-LU. PR1626756

Services Applications

- The `show services l2tp tunnel extensive` and the `show services l2tp session extensive` commands provide incorrect outputs on LTS. PR1601886

Subscriber Access Management

- Subscribers might be stuck in a terminated state when the RADIUS server is unreachable. PR1600655
- The `Service session entry creation failed` errors are seen during ephemeral commit. PR1603030
- The prefix duplication errors might occur for DHCPv6 over PPPoE subscribers. PR1609403
- DHCP session fails with the `session-limit-per-username` statement configuration. PR1612196

- The RADIUS CoA NAK might not be sent with the configured source address in a virtual router environment. PR1625858

User Interface and Configuration

- The `apply-path` configuration does not expand for the configuration under groups. PR1592032
- A low privileged user can elevate their privileges to the ones of the highest privileged J-Web user logged in. PR1593200
- The mustd process might crash with multiple core files due to memory issue. PR1599641
- Invalid JSON and XML output formats appear for the command like `show system resource-monitor ifd-cos-queue-mapping fpc x | display [json|xml]`. PR1605897

VPNs

- The multicast route is not getting installed after exporting of secondary routes from one instance to another. PR1562056
- Unable to add the BGP standard community to NGMVPN Type-6 and Type-7 routes in the VRF export policy. PR1589057
- The backup router rpd process might crash on all Junos OS platforms. PR1594561
- The rpd process might crash if the interface goes down in a BGP-MVPN scenario. PR1597387
- The rpd process might crash during unified ISSU if the `auto-sensing` statement is enabled for l2circuit. PR1626219

Resolved Issues: 21.1R2

IN THIS SECTION

Class of Service (CoS)

- On MPC7E, MPC8E, and MPC9E line cards, the BPS counter of the egress queue displays the wrong BPS value when the cell mode is configured on the static interface. PR1568192
- Unable to configure policer with bandwidth-limit greater than 50g. PR1575049
- Traffic loss might be observed if per-unit-scheduler is configured on AE interface PR1599857

EVPN

- The rpd process might crash under EVPN-VPWS environment. PR1562160
- Prefix added to the mhevpn.evpn.0 output route table triggers TC failure. PR1566429
- ESI preference is not preferred when configured on lo0 for multicast VXLAN. PR1570618
- The multicast traffic loss might be seen in EVPN-VXLAN scenario with CRB multicast snooping. PR1570883

- Configuring static-mac and no-mac-learning simultaneously on the VXLAN interface causes stale MAC/IP entry in the EVPN database. PR1576147
- The mustd.core process generates core file during upgrading or while committing a configuration. PR1577548
- The rpd process might crash if EVPN routing instances or BGP connections flap. PR1581674
- Multicast traffic might loss in EVPN setup with IGMP snooping used. PR1582134
- After device reboot in an EVPN-VXLAN setup with graceful restart, EVPN routes are not advertised to EVPN peers until rpd is up for 180 seconds. PR1586246
- The BUM traffic might lose after triggering NSR in EVPN-MPLS or EVPN-ETREE scenario. PR1586402
- The traffic might be dropped in an EVPN-VXLAN multihomed scenario. PR1590128
- Transit traffic gets dropped post disabling one of the PE-CE link on a remote multihome PE in an EVPN-MPLS A-A setup with a dynamic-list nexthop is configured. PR1594326

Forwarding and Sampling

- After routing restarts, the remote mask that the routing daemon sends might be different from the existing remote mask that the Layer 2 learning daemon had before restart. PR1452990
- User-defined ARP policer is not applied on aggregated Ethernet interface until firewall process is restarted. PR1528403
- In the VXLAN scenario, the locally originated packets have UDP source port 0. PR1571970
- The pfed memory leak might be observed. PR1573285
- The l2ald process might crash on changing the routing instance. PR1584737

General Routing

- The DHCP ALQ is not working as expected. PR1578543
- On MX10003 platforms, despite of having all, AC low PEM alarm is raised. PR1315577
- The SSL-FP logging for non-SNI session. PR1442391
- On MX204 platforms, incorrect log message for PIC1 when changing the configuration from PIC mode to port mode. PR1500429

- Sometimes external 1 pps cTE is slightly above class B requirement of the ITU-T G.8273.2 specification. PR1514066
- The aggregated Ethernet interface might not come up with LFM configured after reboot. PR1526283
- Removing superflous XML tags within syslog strings. PR1528116
- Kernel crash might occur after NSSU while performing GRES. PR1533874
- The dcpfe process might crash and causes FPC to restart due to the traffic burst. PR1534340
- The CFM sessions go down during FRU upgrade stage of ISSU in MX Virtual Chassis. PR1534628
- The spcd process might crash during early initialization. PR1535536
- Sessions creation rate is set to minimal rate after IDS and CPU throttling in place during DDOS attack. PR1544489
- The kmd process might crash when the interface flaps. PR1544800
- FPC might not boot up on MX960 routers in a certain condition. PR1545838
- The performance of Packet Forwarding Engine process on MX204 platforms might be degraded. PR1545989
- The 40 G or 100 G interfaces might flap during unified ISSU if PTP is deactivated on the interfaces on MX Series platforms. PR1546704
- The PTP protocol might get stuck at Initializing state on MX Series platforms. PR1547423
- `HEAP malloc(0) detected!` errors might be seen when adaptive load-balancing is configured on a LAG. PR1547240
- SPC3 might not come up after the system reboot. PR1555904
- FPC crash might occur after flapping the multicast traffic. PR1548972
- When the MX Series device is in the SAEGW-U mode, in rare cases of a double back-to-back failover involving GRES and node association release, some access peers might not be freed even after the sessions count associated with that peer reaches zero. PR1549689
- Deleting or deactivating the PS interface should not be allowed when use by BBE subscriber. PR1550915
- Silent compact flash (/dev/ada1) failure might occur during reboot or start up of the router. PR1551171
- The interface might not come up with 1G optics. PR1554098

- Unified ISSU upgrade might cause a few interfaces to go down. PR1554099
- `Cattle-Prod Daemon received unknown trigger (type Semaphore, id 1)` error messages are seen on the VTY when we issue CLI commands to fetch host route scale. PR1554140
- `CoS WRED Curve: Create Expr Curve: No curve data points!!` error messages are seen when interpolate is configured under drop profile. PR1554220
- The subscriber sessions might be missed but stay in the authd after performing unified ISSU. PR1554539
- The chassisd process might crash with repeated configuration commits on MX204 and MX10003 routers. PR1555271
- The subscriber's RADIUS interim accounting statistics update might not work in some scenario. PR1555492
- FPC with power related faults might get on-lined again once fabric healing has off-lined the FPC. PR1556558
- The dcpfe process might crash and restart with a dcpfe core file created while running the Type 5 EVPN-VXLAN with 2000 VLANs. PR1556561
- The framed route installed for a demux interface has no MAC address. PR1556980
- Script fails while committing the IPsec authentication configuration as the algorithm statement is missing. PR1557216
- Multiple FPCs might crash when performing GRES or FPC reboot repeatedly in a subscriber scenario. PR1557294
- The l3static license is required though it is included in base license. PR1557631
- When 100 G or 40 G interface are configured with protocol PTP, packets corruption is seen. PR1557758
- The MAC addresses learned in a Virtual Chassis might fail due to aging out in the MAC scaling environment. PR1558128
- Application identity unknown packet capture utility does not function when enhanced-services mode is enabled. PR1558812
- Some transmitting packets might get dropped due to the `disable-pfe` action is not invoked when the fabric self-ping failure is detected. PR1558899
- The device might run out of service after GRES or unified ISSU. PR1558958

- The subscriber management infrastructure daemon (smid) process might be stuck at 100 percent. PR1559402
- Single rate three color policer does not work. PR1559665
- Zero suppression disable for for streaming telemetry. PR1559882
- The untagged traffic routed over `native-vlan` might be dropped. PR1560038
- The PTP master line card servo might stuck in Freerun state. PR1560074
- The VXLAN queue DDoS violation and RARP packets flood might happen if receiving the RARP packets more than the supported DDoS bandwidth. PR1560243
- The PIC in SRX5K-SPC3 and MX-SPC3 card might get stuck in offline status after flowd crash occurs on it. PR1560305
- Telemetry might not work after reboot or upgrade. PR1560496
- Interface cannot send or receive packets after repeated link flaps on MPC10 and MPC11E line cards. PR1560772
- The tunable optics SFP+-10G-T-DWDM-ZR does not work. PR1561181
- SPC3 is not supported on MX in 21.1R1 and 20.4R2 for deployment. PR1561188
- After recovering from restart routing immediately, object-info anomalies is observed on rpd agent. PR1561812
- The dcpfe process might crash after deleting VXLAN configuration. PR1562692
- The rpd process might crash when the routing-instances are deleted and recreated quickly. PR1562905
- Layer 2 interface information is not included in DHCPv4 option-82 circuit ID or remote ID DHCPv6 and relay-agent-interface-id or relay-agent-remote-id options when service provider style configuration for switch interface is employed. PR1564010
- Commit error observed when tunnel-service is configured on a PIC without explicit bandwidth. PR1565034
- On MX2010 or MX2020 routers, the following error message might be observed after switchover with GRES and NSR: `CHASSISD_IPC_FLUSH_ERROR`. PR1565223
- Unable to bring up more than one client on one VLAN at the same time. PR1565249
- The KRT log file might continue to grow after removing the KRT log configuration. PR1565425
- The mspmand process crash might be seen on the PIC of MS-MPC/MS-MIC. PR1566325

- Pushing more than 2 MPLS labels might not work. PR1566828
- The rpd process generates core file at boot time of a device. PR1567043
- The chassisd process crash might be seen on MX Series platforms. PR1567479
- TLB composite next hop is installed incorrectly in other routing-instances. PR1567568
- MAC addresses might not be relearned successfully after MAC address age timeout. PR1567723
- The active DHCP subscribers might not get synchronized to backup BNG. PR1567735
- On MX204 routers, FPC might display high CPU utilization because of the JGCI background thread that runs for a long period. PR1567797
- State is not established for `show bgp bmp` *`station name`* post authentication-key bmp-auth configuration. PR1568046
- BFD flaps might be seen between leaf and core during spine reboot causing other protocols flap. PR1568615
- SPC3 card interfaces are not created. PR1568694
- The nsd might crash after turning off the address translation for the NAT rules in the USF scenario. PR1568997
- The rpd process might crash while using BFD API to bring up the BFD sessions. PR1569040
- Traffic loss might be observed when SCU accounting is configured and logical-systems is enabled. PR1569047
- LLDP out-of-bounds read vulnerability in l2cpd. PR1569312
- Wi-Fi mPIM is reaching out to NTP and DNS servers. PR1569680
- The MPLS traffic passed through the back-to-back PE router topology might match the wrong CoS queue. PR1569715
- The mspmand process might crash if the packet flow control issue occurs on MS-MPC or MS-MIC. PR1569894
- The log message `/tmp//mpci_info: No such file or directory :error[1]` might be seen on VM host platform. PR1570135
- The jinsightd process might be stuck with high CPU process utilization. PR1570526
- The bbe-smgd process might crash after committing several thousand addresses in a filter term. PR1570536

- Improve handling deletion of static demux interface with active subscribers. PR1570739
- PDB pull or synchronization does not occur in new primary during unified ISSU. PR1570841
- Upgrading with unified SSDs (2x32G SSD) might result in boot loop in certain scenario. PR1571275
- The toe_gld_toe0_ucode process generates core files at prds_rt_ifl_ipv6_del_hndl_from_desc_list. PR1571279
- Packet loss might be observed when sample based action is used in firewall filter. PR1571399
- FPC crash might be seen when deleting a lot of multicast groups at the same time. PR1571890
- gRPC session hanging in closed state. PR1571999
- The grpcd process might crash and telemetry subscription will retry until grpcd restarts. PR1572107
- DCI traffic loss of 100 percent is observed in transit spine devices. PR1572238
- The TFEB or FPC might fail to come online after rebooting the system or the FPC if interface-set is configured for CoS. PR1572348
- Segment routing might not work properly in IS-IS multiple levels setup. PR1572391
- The `show services mobile-edge sessions summary access-network-peers` command displays incorrect established subscriber output after the UPF handover ENB step. PR1572520
- On MX960 routers, require a fan tray upgrade alarm is raised when the top fan tray 0 is removed, even though the enhanced fan tray is already used. PR1572778
- A traffic loop might be observed after the VCP interface flap. PR1573047
- CFP unplugged message is not logged in Junos OS Release 17.3 and later. PR1573209
- Fabric errors are observed and FPC processes might get offline when MPC3-NG or MPC3E cards are installed along with MPC7/MPC10 and SCBE3/SCB4 operating in increased bandwidth fabric mode. PR1573360
- Some MPC4E-3D line card shows `si5374 clock PLL lock timed out` error message at boot up. PR1573729
- ARP traffic exceeding the policer limit is not discarded. PR1573956
- Only root user is allowed to execute commands on host using vhclient. PR1574240
- QSFP 4x10G interface might not come up after FPC reboot. PR1574279
- DS-Lite throughput degradation might be seen on MS-MPC. PR1574321

- Slow FPC heap memory leak might be triggered by flapping the subscribers terminated over multiple pseudowires. PR1574383
- On the EA-based cards IGMP group membership is displayed incorrectly. PR1575031
- PTP might be stuck in phase acquiring state after unified ISSU. PR1575055
- The rpd process might continuously crash if deleting forwarding-class policy with discard action. PR1575177
- The MPC10E line cards generates the following error message: `user.err aftd-trio: [Error] Em: root: Insert entry failed, entry:parentToken:747441 entryMask:ffffffffffffffff index:52`. PR1575310
- On MX150 routers, the interface might take a long time to power down while rebooting, powering-off, halting, or upgrading. PR1575328
- The `show services service-sets statistics syslog` command returns the following error message as the service-set does not have the syslog configuration: `error: usp_ipc_client_recv_ 1237: ipc_pipe_read fails! error:No error: 0(0), tries:1`. PR1576044
- IPsec tunnel is not established when receiving the proxy-id list. PR1576071
- On MX10016 routers, when the fan X failed alarm is cleared in the fan tray 1, the fan/blower OK SNMP traps are generated for the fan tray 0 [Fan 31 - 41] and fan tray 1 [Fan 11 - 41]. PR1576521
- The LLDP neighbor information displays hex string instead of chassis ID when subtype 1 is used. PR1576721
- The MS-MPC and SPC3 might reset on receiving the subscriber traffic. PR1576946
- The following commit failure error is observed: `Modified IFD "ae0" is in use by targeted BBE subscriber, commit denied - mtu config changed (1522), (1514)`. PR1577007
- The OSPF session over IRB might not come up in the EVPN-VXLAN scenario. PR1577183
- Traffic loss might be seen when subscriber service over aggregated Ethernet bundle interface. PR1577289
- When line card is booted on Routing Engine 1 being master, nextgen statistics failed to fetch the value of backup MAC address correctly. PR1577611
- Native sensors does not work for LDP LSP. LDP P2MP sensor. PR1577931
- The bbe-smgd process crash might be seen when the RADIUS server sends multiple CoA. PR1578162
- TACACS traffic might be dropped. PR1578579
- High FPC CPU usage might be seen when signal on the link is unstable. PR1579173

- Random or silent reboot might be seen. PR1579576
- On the MPC11E line cards, system resource monitor does not list some of the available Packet Forwarding Engines. PR1579975
- Authentication might fail if the password contains special characters. PR1580003
- ON MX Virtual Chassis, gRPC-based components or sensor output is missing a lot of data. PR1580120
- While mapping analyzers to the channelized port, mirror might not work properly. PR1580473
- The l2cpd process might crash on platforms with dual Routing Engines. PR1580479
- More than one subscriber on same vlan fails to apply same FWF template. PR1580826
- Need to add support for Virtual Chassis licensing. PR1580880
- The following error message is observed: `kern.ipc.maxpipekva exceeded; see tuning error`. PR1581192
- Memory leak might happen due to stale NAT64 entries. PR1581231
- The rpd process might crash on the new primary Routing Engine after performing graceful switchover. PR1581878
- Changing the bandwidth statement does not take affect for SNMP ifHigSpeed oid until a PSX interface is disabled and enabled. PR1582060
- The voice VLAN might not get assigned to the access interface. PR1582115
- Communication between two CEs might be failed when BGP rib-sharding is enabled. PR1582210
- The rpd process might generate core file after Routing Engine switchover. PR1582095
- The rpd process might stuck in 100 percent due to race condition. PR1582226
- The bbe-smgd crash might be seen after subscriber log out due to a rare timing issue on MX platforms. PR1582356
- On MX960 devices, the 400 G and 4x100 G optics laser restores after reboot despite interface disable. being configured. PR1582418
- Destination port might be incorrectly set on MS-MPC/MS-MIC in DS-Lite scenario. PR1582595
- Node locked license addition fails. PR1582704
- Configuring or removing `hierarchical-scheduler` or `per-unit-scheduler` might cause traffic to stop forwarding. PR1582724
- The firewall filter logs are incorrectly populated the protocol 8847 entries. PR1582780

- Reset JBS, JAS, and JPS definition to align with new license model. PR1583438
- Reset PFL and AFL definition to align with new license model. PR1583439
- The firewall filter cannot be programmed after deleting a large filter and adding a new one in a single commit. PR1583440
- SNMP SysObjectID.0 is empty with unified-services enabled. PR1583534
- TCP connection to syslog server might fail to be established after adding the `tcp-log` configuration for an existing service-set. PR1583979
- Layer 2 multicast VXLAN instance is down since local VTEP logical interface is not associated to EVPN instance. PR1584109
- The jsd process hogging CPU. PR1584357
- Traffic might not get filtered properly when security intelligence profile is configured on the MX platforms. PR1584377
- After performing NSSU, `timeout waiting for response from fpc0` error message is seen while checking version detail. PR1584457
- The node name must not be attached to the system hostname under LLDP. PR1593991
- The rpd process might crash due to a rare timing issue if both BGP Local-RIB and Adjacency-RIB-In route monitoring are enabled in BMP. PR1584560
- Bridge domain names information is not displayed properly for the `show bridge statistics instance` command. PR1584874
- After changing configuration, the `show bridge statistics` shows extreme large value. PR1584876
- Traffic impact might be seen when tunnel-services bandwidth is configured. PR1584969
- GRE OAM packets are sent through queue 0 with the force-control-packets-on-transit-path statement enabled. PR1586169
- Traffic drop is seen after enabling flexible-queuing-mode on MPC2E line cards. PR1586403
- The l2ald process might crash on changing the routing instance. PR1586516
- Inter and intra VNI traffic drop might occur in spine with EVPN-VXLAN CRB configuration. PR1586537
- The rpd core file might be observed if executing the `show igmp continuous stats` command after GRES. PR1587023

- The SNMP trap for MAC notifications might not be generated when an interface is added explicitly under switch-options. PR1587610
- The bbe-smgd process might crash if the staled ACI based subscribers are not cleaned up properly. PR1587792
- The rpd process crash might be observed on the router running in a scaled setup. PR1588439
- The bbe-statsd memory leak might be observed on backup Routing Engine during subscribers login or logout. PR1589081
- The jsd process might crash in a rare condition in a telemetry scenario. PR1589103
- The l2cpd process might crash. PR1589216
- Allow default license for FBF, CFM, VRRP, QINQ, MC_LAG, TIMING, IGMP, PIM, GRE_TUNNEL, RIP, OSPF, Virtual Chassis, and sFlow. PR1589920
- Traffic loss might be observed for interface configured in subnet 16. PR1590040
- VXLAN DDoS violation might occur when disabling the port mirror analyzer output interface. PR1590150
- Traffic loss might be observed due to FPC crash in a scaled subscriber scenario. PR1590374
- Non-zero values might be displayed against the drop field in the `show network-agent statistics` command output post switchover scenarios. PR1590432
- NAT service might not happen after performing AMS switchover or deactivating and activating NAT service. PR1590890
- Traffic loss might be observed after changing SAK keys. PR1591432
- If the CoS CR-features used by VBF service is configured, MPC might crash with subscriber. PR1591533
- Frequent phydriver sync_state toggling results in high 2-way time Errors. PR1591667
- If the CoS CR-features used by VBF service is configured, MPC might crash with subscriber. PR1591533
- On all Junos platforms, xSTP might not get configured when enabled on an interface with SP style configuration. PR1592264
- The aftmand process might crash when an interface is configured with the analyzer. PR1592267
- The mobiled daemon might crash after switchover is performed for an AMS interface or crash occurs on service PIC where the AMS member interfaces are present. PR1592345

- Routing Engine kernel might crash due to logical interface of aggregated interface adding failure in Junos kernel. PR1592456
- Using the BITS interface from backup Routing Engine for clock recovery might not work. PR1592657
- The packet comes from the PS interface and forwards to the SPC3 might be dropped. PR1592706
- Any mmcq based services might crash due to shared memory queues issue happens in a rare condition. PR1592889
- The TCP keepalive might not be processed by the private network host. PR1593226
- Fabric errors will be generated after swapping MPC10E line card with MPC7E line card in the same slot. PR1593821
- Packet drop might be seen when traffic is moving from one FPC to another FPC. PR1594244
- On MX5, MX40, and MX80 routers, TEB stuck in present state. PR1595107
- Platforms with EVPN-VXLAN with shared-tunnel configuration, when there is BGP flap or restart of l2ald, info logs appear in VTY. PR1595203
- Firmware might fail to be downloaded to MIC on MX Virtual Chassis setup. PR1595693
- The l2ald process might crash on all leaves and spines after a new leaf is added to the EVPN fabric. PR1596229
- CGNAT MX SPC3 AMS warm-standby 1:1 redundancy problem with CLI, CPU statistics lost data after PIC failover. PR1596976
- Major alarms on all FPCs in chassis after some time from bootup. PR1597066
- Subscriber management daemons might continuously core and shutdown with Routing Engine sensors invalid configuration. PR1598351
- The AFEB crash might be observed with MIC-3D-8DS3-E3. PR1598411
- Component sensor is not exporting logs for /components/component[name='Chassis']/state/description. PR1598816
- MX SPC3 applications for protocol ICMP is not detected and does not allow user to modify inactivity-timeout values. PR1599603
- On an MX Series router with MPC3E (non 3D) and SCB3E, a fabric plane might go in to check state after configuring increased bandwidth mode. PR1602080
- The Packet Forwarding Engine might be disabled by a detected major CMERROR event while ungracefully removing the MIC from MPC2E-3D-NG and MPC3E--3D-NG. PR1602939

- In a subscriber management scenario, under a rare condition, the Routing Engine reboots and generates vmcore files. PR1607282

Infrastructure

- A debug log generated when SDK code invokes malloc with the size as 0 while destroying a multicast entry. PR1546036
- The net installation (PXE) is not working. PR1577562
- Some MAC addresses might not be aged out. PR1579293

Interfaces and Chassis

- On the MPC10 line cards, DMRs or SLRs are not received with an EVPN up MEP on the aggregated Ethernet interface with normalization. PR1543641
- Block duplicate IP across different logical interfaces inside same routing instance. PR1555861
- MAC address entry issue might be observed after the MC-LAG interface. PR1562535
- Unable to set member-id as Routing Engine is in synching mode forever when its having invalid Virtual Chassis data. PR1569556
- On MX Series routers, if-media-type is missing from interface XML output. PR1574035
- There might be increase in memory for the fabspoked process. PR1574391
- MX Virtual Chassis ISSU incompatible FRU offline can result in unexpected FPC restarts after unified ISSU completes. PR1575687
- Error messages are seen during GRES. PR1575689
- MC-LAG interfaces might go down if the same VRRP group-id is configured on multiple IRB units. PR1575779
- ARP resolution failure might occur during VRRP failover. PR1578126
- Newly added MC-LAGs do not come up after Routing Engine switchover. PR1583547
- Add configuration for PPP NCP max-failure number of retry count. PR1584168
- Unable to configure pseudowire interface on an MX10003 in virtual chassis mode. PR1587499
- The dcd process crash might be seen after performing Routing Engine switchover, reboot, or management interface configuration change. PR1587552

- On MX240 platforms, difference between the statistics of DMM sent and DMR received is not as expected. PR1595780
- The VRRP host cannot be reached if `native-vlan-id` is configured. PR1595896

Intrusion Detection and Prevention (IDP)

- Adding signature in packet drop reason and sending to record packet drops module. PR1574603

J-Web

- J-Web allows a locally authenticated attacker to escalate their privileges to root. PR1511853
- To improve performance in Monitoring > Network > Interfaces page, admin status is removed, services and protocols data merged into one host inbound traffic. PR1574895

Juniper Extension Toolkit (JET)

- The custom JET APP will be lost after rebooting. PR1570563

Layer 2 Features

- LACP does not come up in the non-oversubscribed mode for a set of ports. PR1563171
- Traffic forwarding for VLAN 2 might not be correct when a VLAN member is removed from the ESI interface. PR1570446

Layer 2 Ethernet Services

- Copying of files to the RCB over WAN ports is slow. PR1496895
- Aggregate Ethernet interface flap might be seen during NSSU. PR1551925
- DHCP packet drop might be observed when the DHCP relay is configured on a leaf device. PR1554992
- The option 82 information is incorrectly cleared by the DHCP relay agent. PR1568344
- The DHCP client will be offline for 120 seconds after sending the DHCPINFORM message in the DHCP relay scenario. PR1575740
- The DHCP relay drops packets during the renewal DHCP process. PR1576417
- The jdhcpd process might crash if relay-source lo0 is enabled in DHCP relay. PR1580724

- There is ALQ synchronization issue on priary BNG and backup BNG with loss of subscriber session redundancy via PS interface. PR1583310
- The DHCP ALQ queue might get stuck causing subscriber flap. PR1590421
- The jdhcpd process might not respond to any discover message when it is in `clients waiting to be restored` state. PR1592552
- The jdhcpd core file post Junos OS upgrade. PR1594371

MPLS

- The rpd process might crash in corouted bidirectional RSVP LSP scenario. PR1544890
- Incorrect EXP bit change might be seen in certain conditions under MPLS scenario. PR1555797
- Traffic loss might be observed during rpd process crash when RSVP signaled P2MP LSP is configured. PR1559022
- Traffic sent over an LSP might be dropped if two consecutive PLRs along the LSP perform local repair and bypass protecting the second PLR fails. PR1566101
- Unexpected LSP packet count is observed in the ingress MPLS LSP statistics. PR1570382
- The rpd process on the transit node might crash when MPLS traceroute on the ingress node is performed. PR1573517
- The rpd process generates core file when deactivating PCEP protocol followed by RSVP protocol. PR1579370
- Sub-optimal routing issues might be seen in case LDP route with multiple next hops. PR1582037
- Add the `lsp-ping-multiplier` option for LDP-OAM similar to RSVP-OAM. PR1582254
- The LDP replication session might not get synchronized when the dual-transport statement is enabled. PR1598174
- VPLS connection might get down if dual-transport statement is configured. PR1601854
- VPLS connection might get down if the dual-transport statement is configured. PR1601854

Multicast

- Multicast traffic in MVPN setup might be silently dropped and discarded on platforms acting as transit LSR. PR1555274
- FPC might crash in a multicast scenario. PR1569957

Network Address Translation (NAT)

- Services NAT mappings and sessions are incorrect while checking the SIP sessions from public to private and RTP from private to public. PR1577922

Network Management and Monitoring

- SSH connection might become unresponsive and logs show `kern.maxfiles limit exceeded by uid` messages. PR1567634
- Slow memory leak could be observed for snmpd process. PR1575790

Platform and Infrastructure

- Traffic loss might be observed due to FPC crash on MX Series platforms. PR1482683
- Interwork failure between Junos OS as RPM client and TVP platforms as RPM server and vice versa. PR1508127
- Console access on backup Virtual Chassis member is not allowed. PR1530106
- The npc process generates core file in `igmp_process_wakeup_events,igmp_pfe_thread,thread_detach_tty`. PR1534542
- CoS queue egress interface forwarding-class might not work as expected. PR1538286
- The following major error message might cause the Packet Forwarding Engine to disable: `XQ_CMERROR_SCHED_L3_PERR_ERR`. PR1538960
- Subscribers over an interface-set might not be able to log in. PR1539260
- Upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart. PR1557881
- The BUM frame might be duplicated on an aggregate device if the extended-port on the satellite device is an aggregated Ethernet interface. PR1560788
- Multicast traffic with incorrect source MAC address might be observed from IRB interface. PR1561313
- Traffic loss might be observed due to FPC crash on MX Series platforms. PR1563144
- The `enforce-strict-scale-limit-license` configuration enforces subscriber license incorrectly in the ESSM subscriber scenario. PR1563975

- The `Last flapped` timestamp for interface fxp0 gets reset every time the `monitor traffic interface fxp0` command is executed. PR1564323
- PFEX might crash when soft error recovery feature is enabled on the Packet Forwarding Engine. PR1567515
- The L2TP tunnel might not work with filter-based encapsulation. PR1568324
- On MX platforms, toe_lu_stats_ucode core is found at `jbeta_fcv_alloc_fcv_idx_global jbeta_sfilter_fcv_cb bwy_dfw_sfilter_fcv_cb`. PR1569328
- On MX Series with MPCs/MICs, subscribers error logs might be seen. PR1570631
- FPCs might crash randomly while deleting the interface-set in the system. PR1571192
- On platforms with EVPN-VXLAN configured, the next hop memory leak in MX Series ASIC happens whenever there is a route churn for remote MAC-IP entries learned bound to the IRB interface in EVPN-VXLAN routing instance. When the ASIC's next hop memory partition is exhausted, the FPC might reboot. PR1571439
- Scale-subscriber license might not be updated properly on the backup Routing Engine which leads to `License grace period for feature scale-subscriber(44) is about to expire` alarm after GRES. PR1573289
- Uninitialized Read Error at EDMEM[0x7cb601b0]. PR1573920
- Introduce two new major CMERRORs for XM chip-based line card to stabilize the running device. PR1574631
- Memory partitioning issue might happen on Packet Forwarding Engine after applying sampling and flex-flow-sizing to the MX Series with MPCs/MICs based line-cards. PR1575994
- If committing `source-address` *`addr`* `routing-instance` and then delete `source-address` *`addr`* in private edit mode, commit fails with warning message. PR1582529
- VRRP device originally taking slave role might cause destination IP unreachable after VRRP mastership switchover. PR1584115
- FPC might crash in a scaled-firewall configuration. PR1586817
- The traffic might not failover with shared-bandwidth-policer enabled on aggregated Ethernet PR1588708
- Upon receipt of specific sequences of genuine packets destined to the device, the kernel will crash and restart. PR1595649
- VLAN tagged traffic might be dropped with service provider style configuration. PR1598251

- The service filter might get incorrectly programmed in acket Forwarding Engine due to a rare timing issue in enhanced subscriber management environment. PR1598830
- The kernel core file might be seen if restarting BGP connections after deleting BGP authentication. PR1601492

Routing Policy and Firewall Filters

- The dns-name cannot be resolved if customer-defined routing instance is configured under name-server. PR1539980
- The rpd process might crash when the deletion of routing table occurs. PR1565629
- The rpd process might crash due to the source-address-filter-list statement enabled within the policy. PR1565891
- Traffic loss might be observed due to rpd process crash when NSR switchover is performed. PR1579830
- On MX Series platforms, bbe-smgd - dymanic-profile NACK due to configuration error reading address mask prefix-length in policy-options or policy-statement. PR1583535

Routing Protocols

- Traffic might be silently discarded when a BGP route that is part of multipath gets deleted. PR1514966
- Route validation states might flip between VALID, INVALID, and UNKNOWN in some corner case. PR1556656
- The ISO routes are not leaked in default (master) instance after switchover or reconfiguration. PR1558532
- When admin color based policy evaluation happens with the policy LFA configuration, the backup next hop chosen (among the different backup next hops possible) might not be correct. PR1558581
- Incorrect Active, Received, and Accepted counters might be seen in the output of the `show bgp summary`. PR1558678
- The ppmd memory leak might cause traffic loss. PR1561850
- There might be traffic loss when GRE interface flaps. PR1566428
- The rpd process might crash in BGP L2VPN scenario due to memory corruption. PR1567026

- The rpd memory leak might be observed during CLI or ephemeral commits in OSPFv2 scenario. PR1568157
- The rpd process might crash continuously when MoFRR is configured along with TI-LFA. PR1568750
- Traffic might be lost during mirror data transmit from the primary ppmd or bfdd. PR1570228
- There might be 10 seconds delay to upload the LSP on the point-to-point interface if rpd is restarted on its direct neighbor. PR1571395
- SNMP MIB ospfv3NbrState is returning drifted value. PR1571473
- After the first parallel ISSU, subsequent ISSU aborts with `Aborting Daemon Prepare`. PR1572265
- The BFD session of DHCP subscriber does not come up on the MPC2E line card and gets stuck in the `Down` state. PR1572577
- The DHCP packets might get drop in the static VXLAN scenario. PR1576168
- The ppmd might crash when enabling MD5 authentication on OSPF with BFD flapping. PR1576893
- BGP session flap might be observed after the Routing Engine switchovers when the VRRP virtual address is used as the local address for the BGP session. PR1576959
- Multicast traffic loss might be observed due to logical PIM de-encapsulation, interface is not created as expected. PR1577461
- The rpd process might crash when two or more routing instances are deleted in one shot. PR1578740
- The dcpfe process might crash when any interface flaps. PR1579736
- Traffic loss might be observed due to rpd process crash when tunnel encapsulation is used. PR1579818
- The BGP session carrying VPNv4 prefix with IPv6 next hop might be dropped. PR1580578
- BGP replication might be stuck in rare and timing conditions. PR1581578
- The rpd process might crash in BGP and MPLS scenario. PR1581794
- Route resolution issue after controller facing Packet Forwarding Engine restart or core interface disable and enable. PR1581845
- Possible rpd process crash with the `routing-options transport-class` configuration during the routing restart. PR1582081
- With IGMP snooping implemented, there is unexpected jitter issue that could cause traffic loss. PR1583207

- The rpd process crash might be seen in certain IS-IS scenario. PR1583484
- On rare occasion, rpd process generates core file on backup Routing Engine after loading a new image. PR1583630
- Origin-validation (RV) replication status shows up in the `show task replication` even when not configured. PR1583692
- The rpd process might crash after committing with the configured static group. PR1586631
- Incorrect BGP next hop advertisement in a L3VPN scenario. PR1587879
- Multicast traffic loss could be observed after unified ISSU is being performed. PR1588555
- The rpd process might crash in a scaled routing instances scenario. PR1590638
- PIM joins might not be synchronized between master and backup Routing Engines because of ppmd restart. PR1591685
- Doing BGP disable or enable in a short time interval on a scaled NSR router can result in backup rpd restart. PR1591717
- The rpd process crash might be seen if BGP peer flaps. PR1592123
- The remote LFA backup path might not be formed. PR1592424
- BGP egress-TE routes lose to BGP routes using the same protocol preference. PR1593332
- The routing process might crash due to memory corruption while processing BGP multipath route. PR1594626
- The rpd process might be stuck at 100 percent in OSPFv3 scenario. PR1601187
- With the rib-sharding statement enabled, any commit will flap all BGP sessions with 4 byte peer-as (AS number 65536 or greater). PR1607777

Services Applications

- The CoA with LI-on or LI-off message might be dropped during CoA process. PR1554618
- IWF AVP value might not be reflected properly on LTS. PR1581096
- The `show services l2tp tunnel extensive`, `show services l2tp session extensive`, and `show subscribers accounting-statistics` commands do not work on LTS. PR1596972

Subscriber Access Management

- BBE-SMGD configures incorrect vbf_accurate_accounting_bits to the Packet Forwarding Engine. PR1515899
- The authd process might crash after performing unified ISSU in a MX BNG scenario. PR1570096
- CoA request might not be processed correctly from time to time. PR1571501

User Interface and Configuration

- The mustd process might crash with multiple cores files due to memory issue. PR1599641

Virtual Chassis

- Virtual Chassis port might not come up after upgrade when QSFP+-40G-SR4, QSFP+-40G-LR4, or QSFP+40GE-LX4 is used PR1579430

VPNs

- Traffic from the reverse direction might cause traffic loss for up to 1 second with NSR switchover. PR1558395
- The rpd process might crash during a race condition under BGP multipath scenario. PR1567918
- The iked process might crash when IKEv2 negotiation fails on MX Series devices. PR1577484
- The rpd process might crash in the NG-MVPN scenario. PR1579963
- The traffic of the draft-rosen multicast VPN might lose after switching over the Routing Engines. PR1584720
- Unable to add BGP standard community to NGMVPN Type-6 and Type-7 routes in VRF export policy. PR1589057
- The ddos-protection reason `packets failed the multicast RPF check` might be seen in NG-MVPN scenario with GRE transport. PR1591228

Resolved Issues: 21.1R1

IN THIS SECTION

Class of Service (CoS)

- While configuring the WRED profile to a scheduler, you can use either any/any not-any/not-any combination of protocol or loss priority. PR1524259
- The explicit classifier or rewrite-rule might not work as expected for a logical interface if the wildcard configuration is also applied. PR1556103

EVPN

- `no-arp-suppression` is required for MAC learning across the EVPN domain on the static VTEP. PR1517591
- ARP replies from the CE device gets dropped incorrectly at the PE fdevice or the EVPN routes resolving through MPLS-over-UDP. PR1563802

  The l2ald process might crash under the VLAN-based EVPN-VxLAN scenario. PR1550109
- The BUM traffic might get dropped in the EVPN-VXLAN setup. PR1525888

- The route table shows additional paths for the same EVPN or VXLAN Type 5 destination after upgrading from Junos OS Release 18.4R2-S3 to 19.4R1-S2. PR1534021
- All the ARP reply packets toward some address are flooded across the entire fabric. PR1535515
- The GE LOS alarm logs on the change in `IFF_CCCDOWN` are not logged in the syslog message file. PR1539146
- The rpd memory might leak when the EVPN configuration is changed. PR1540788
- The l2ald process might generate a core file when the EVPN-VXLAN configuration is changed. PR1541904
- The rpd might crash after adding route-target on a dual-Routing Engine system under the EVPN multihoming scenario. PR1546992
- VLAN ID information is missed while installing the EVPN route from the BGP Type 2 Route after modifying a routing instance from instance type EVPN to instance type virtual-switch. PR1547275
- Remote code execution vulnerability is observed in the Overlayed service. PR1517591

Forwarding and Sampling

- The srrd process might crash in a high route churns scenario or if the process flaps. PR1517646
- The l2ald process might crash when a device configuration flaps frequently. PR1529706
- VLAN-ID-based firewall match conditions might not work for the VPLS service. PR1542092
- MAC learning issue might occur when EVPN-VXLAN is enabled. PR1546631
- All traffic is dropped on the aggregated Ethernet bundle without VLAN configuration if bandwidth-percent policer is configured. PR1547184
- The l2ald process might crash due to a next-hop issue in the EVPN-MPLS. PR1548124
- Configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. PR1563641

General Routing

- Dynamic tunnel summary displays a wrong count of up and total tunnels. PR1429949
- The riot might crash due to a rare issue if vMX run in the performance mode. PR1534145
- The BFD sessions might not come up in the VXLAN scenario. PR1538600
- Unable to show to which shard a given route is hashed. PR1430460

- On the MPC11E line card, the number-of-sub-ports configuration on the 4x10GbE channelized ports might cause the channels to go down. PR1442439
- The MPC2E-NG or MPC3E-NG card with a specific MIC might crash after a high rate of interface flaps. PR1463859
- The following line-card errors are seen:

```
HALP-trinity_nh_dynamic_mcast_add_irb_topo:3520 snooping-error: invlaid IRB topo/ IRB ifl
zero in l2 nh 40495 add IRB.
```

  PR1472222
- Dynamic SR-TE tunnels do not get automatically re-created at the new primary Routing Engine after the Routing Engine switchovers. PR1474397
- Memory utilization enhancement is needed. PR1481151
- Subscribing to `/linecard/packet/usage` and triggering the UDP decoder, the hardware statistics are exported with improper hierarchy. PR1485739
- Prefix is not emitted for the `te-lsp-timers/state/cleanup-delay` sensor path for OCST. PR1500690
- Transit IPv4 traffic forwarding over BGP SR-TE might not work. PR1505592
- The log file to log the activities associated with the `request rift package activate` command is created with the permissions of the user. If multiple users run the command, the command might fail due to the write permission error. PR1514046
- On the MX960 routers, the `show interfaces redundancy rlt0` statement shows current status as **Primary down** as the FPC is still in the **Ready** state after RLT failover (restart FPC). PR1518543
- The BFD session status remains down at the non-anchor FPC even though BFD session is up after the anchor FPC reboots or panic. PR1523537
- The rpd process might crash when the routing-instances are deleted and recreated quickly. PR1562905
- FPC might not be recognized after the power cycle (hard reboot). PR1540107
- No response from the other Routing Engine for the last 2 seconds triggers the following SNMP trap message:

```
Fru Offline
```

PR1524390

- Problem with static VLAN deletion with active subscribers, and the FPC might be stuck at the **Ready** state during restart. PR1525036
- The following error message is observed during GRES if an IRB interface is configured without a profile:

```
RPD_DYN_CFG_GET_PROF_NAME_FAILED.
```

PR1526481

- The l2cpd process might crash when removing LLDP on an aggregated Ethernet interface. PR1528856
- The `speed` command cannot be configured under the interface hierarchy on an extended port when the MX204 or MX10003 router works as an aggregation device. PR1529028
- The SFP-LX or SFP-SX optics on MIC-3D-20GE-SFP-E/EH might show as unsupported after ISSU. PR1529844
- The following error message for port might be seen:

```
FAILED(-1) read of SFP eeprom
```

PR1529939

- On the MX2010 routers, BiDi 1G SFP optics gives wrong value in JVision for `optics/laser_rx_power_*_thresholds`.

PR1530120

- After performing ISSU with a high-scale bridge-domain configuration, less than 0.0254 percent of traffic loss is observed for a single bridge-domain interface. PR1531051
- On MX204 and MX10003 routers, PEM 0 always shows as absent or empty even if PEM 0 is present. PR1531190
- On the MX150 routers, configuring the `no-flow-control` statement under gigether-options does not work. PR1531983
- Wavelength unlocked alarm is `On` when using SFP+-10G-T-DWDM-ZR optics. PR1532593
- The interface with the `pic-mode 10GE` configuration might not come up if upgraded to Junos OS Release 18.4R3-S4 or later. PR1534281

- Some routes might get incorrectly programmed in the forwarding table in the kernel that are no longer present in rpd. PR1534455
- PTP slave might discard the PTP packets from primary when MPLS explicit-null is configured. PR1547901
- Packets drops might be seen after configuring the PTP transparent clock. PR1530862
- PTP slave might discard the PTP packets from primary when MPLS explicit-null is configured.PR1547901
- The log file of the lcklsyncd process shows empty.PR1567687
- On the ACX710 routers, continuous reboot due to configuration under auxiliary port s observed. PR1580016
- Multiple vmxt processes might generate core files. PR1534641
- The MPLS traffic that passes through the back-to-back PE device topology might match the wrong COS queue. PR1569715
- The following log message might be seen on VM host platform: `/tmp//mpci_info: No such file or directory :error[1]`PR1570135
- SNMP MIB walk for jnxSubscriber OIDs returns a general error message. PR1535754
- All SFBs might go offline due to fabric failure and fabric self-ping probes performing the `disable-pfe` action. PR1535787
- Enhancements are needed to debug l2ald. PR1536530
- The chassisd memory leak might cause traffic loss. PR1537194
- The following error message might be observed when the JAM packages for the MX204, MX10003, and MX10008 routers are installed:

  ```
  AM: Plugin installed for summit_xxx PIC
  ```

  PR1537389
- Version-alias gets missed for subscribers configured with dynamic profiles after ISSU. PR1537512
- Deactivating or activating PTP or SyncE in the upstream router causes the 100GbE links on the LC2103 to flap. PR1538122

- The The MPC10 and MPC11 Packet Forwarding Engine FPCs (MPC10 and MPC11 line cards) Packet Forwarding Engine `show jnh exceptions inst <inst-number>` command might cause the FPC to crash. PR1538138
- Traffic drop might be seen while executing the `request system reboot` command. PR1538252
- The accounting interim-updates for subscriber does not work after GRES and subsequent reboot of FPCs in the node-slicing setup. PR1539474
- The rpd memory might leak on the backup Routing Engine due to link flaps. PR1539601
- The mspmand process leaks memory in relation to the MX Series telemetry, reporting the following error message:

  ```
  RLIMIT_DATA exceed
  ```

  PR1540538
- With hold-time configuration, the ge interfaces remain down on reboot. PR1541382
- Subscriber might not come up on some dynamic VLAN ranges in a subscriber management environment. PR1541796
- The dcpfe process might crash and restart with a dcpfe core file created while running the Type5 EVPN-VXLAN with 2000 VLANs. 1556561
- Packets corruption on 100G or 40G interface is observed when configured with protocol PTP. PR1557758
- During ISSU, BNG losses the subscriber sessions without sending Session Stop but stay in authd. PR1554539
- The l2alm process high CPU utilization might be observed in the EVPN-VxLAN environment. PR1551025
- After changing addresses in the source pool, if the carrier-grade NAT traffic does not stop, the source pool cannot perform the NAT translation from the new pool. PR1542202
- The KRT queue might get stuck after the Routing Engine switchover. PR1542280
- Port mirroring with maximum-packet-length configuration does not work over the GRE interface. PR1542500
- The mspmand process might generate a core file on activating or deactivating the interface. PR1544794

- The riot forwarding daemon might crash on vMX-based platforms configured with an IRB interface. PR1544856
- Traffic loss might be observed when the Switch Fabric Board 3 and MPC8E 3D combination is used in the MX2010 or MX2020 router. PR1544953
- The FPC process might crash during the system booting. PR1545455
- RPD fails to program new routes, and continuous rpd errors might be observed. PR1545463
- Plane offline IPC of chassis-id might time out on MX Series devices with MPC11E line cards. PR1546449
- Unexpected log messages appear related to Neighbor Solicitation (NS) messages with multicast as source address. PR1546501
- Backup Routing Engine vmcore might be seen due to absence of next-hop acknowledgment Infra. PR1547164
- In the syslog output, the `sylog-local-tag` name is truncated as `SYSLOG_SF` when the `sylog-local-tag` name is configured as `SYSLOG_SFW`. PR1547505
- The nsd daemon might crash after configuring inline NAT in USF mode. PR1547647
- `SENSOR APP DWORD` leak is observed during the period of churn for routes bound to the sensor group. PR1547698
- SR-TE might stay up when the routes are deleted through policy. PR1547933
- Multicast traffic drop might be seen after ISSU. PR1548196
- Validation of OCSP certificate might not go through in case of certain CA servers. PR1548268
- Error messages are observed as the backup peer does not send marker acknowledgment for the last 360 seconds for vks 0 slave_ack=0 during ISSU. PR1550492
- The adapted sample rate might be reset to the configured sample rate without changing the sampling rate information in sFlow datagrams after enabling sFlow technology on a new interface. PR1550603
- The rpd might crash when BGP service route is resolved over color-only SR-TE policy. PR1550736
- The PPPoE subscribers might fail to log in. PR1551207
- Slow FPC heap memory might leak due to the flapping of the subscribers terminated over multiple pseudowires. PR1574383
- The Packet Forwarding Engine might get disabled when major CMERROR occurs due to the parity errors. PR1551353

- Disable-pfe with intermittent `ipc_pipe_get_packet(): packet_get()` failed error message and `CM_CMERROR_FABRIC_SELFPING failure` messages are observed. PR1554209
- The following error message might be observed.

```
LCM Peer Absent
```

PR1551760

- Fixed Packet Forwarding Engine instance processing in JnhHandleReplicate to honor the Packet Forwarding Engine mask is observed. PR1553400
- Fabric errors are observed and the FPC processes might go offline with SCBE3, MPC3E-NG, or MPC3E line cards and MPC7 or MPC10 line card in the increased-bandwidth fabric mode. PR1553641
- Configuring HFRR (for example, link-protection) on an interface might cause rpd to crash. PR1555866
- Chassisid SNMP trap `Fru Offline` is not generated on MPC11E line card due to no power. PR1556090
- ISSU might be aborted on the MX Series devices for Junos OS Release 20.2R2-S1. PR1557413
- On the MX150 routers, the following continuous license error is observed:

```
[licinfra_set_usage_nextgen_async:1733] Invalid input parameters.
```

PR1559361

- On the MX960 routers, mismatch between YANG schema and RPC output are observed. PR1559810
- When the system has only one plane (in the process of plane offline or online), the MPC10-10c line card displays a destination error. PR1560053
- The `request system software validate` command might corrupt installation of junos-openconfiguration package. PR1560234
- On the MX240 routers, R0 overlay ping fails. PR1560408
- The l2cpd process might generate a core file on reboot. PR1561235
- On the MX240 routers, the VIA headers do not change properly when the SIP ALG is enabled. PR1561312

- Traffic drop might occur on all platforms running Junos OS when a GRE-based dynamic tunnel is configured. PR1561721
- The rpd might crash during processing huge amount of PIM prune messages. PR1561984
- The following error message might be seen after ISSU:

```
Turbotx process not running
```

  PR1564418
- The PPPoE service-name-tables do not correctly count active sessions matching the agent-specifier aci/ari used for delay. PR1565258
- The MX204 FPC might show high CPU utilization because the JGCI background thread runs for a long period. PR1567797
- On the MX150 routers, the `request system software add` command is disabled in Junos OS Release 19.4R3-S1, 20.1R2, and 20.4R1. PR1568273
- The rpd might crash while using BFD API to bring up BFD sessions. PR1569040
- The agent sensor `__default_fabric_sensor__` seems to be partly applied to some FPCs, which causes the following zero payload issue:

```
AGENTD received empty payload for pfe sensor __default_fabric_sensor__
```

  PR1569167
- GRE OAM keepalive fails to start after the Packet Forwarding Engine reboots. PR1569790
- Fabric errors are observed on a system with MPC3E line cards and MPC4E or MPC5E line card with enhanced MX960 backplane. PR1573360
- DHCP discover packet might be dropped if the DHCP inform packet is received first. PR1542400
- The `show dynamic-profile session client-id` command displays only one IPv6 framed-route information. PR1555476
- On the MX2010 routers, many chassid and fabric related errors are observed after ZPL ISSU. PR1558626
- On the MX480 routers, the MPC10E line cards are restarted after performing GRES with scaled configurations. PR1561259
- On the MX2020 and MX960 routers, the PTP state gets stuck in the `Acquiring` state. PR1562267

- On the MX2010 routers, the aft-ulcd process crashes and generates core files continuously and SLC keeps restarting after upgrade. PR1578191
- On the MX480 routers, the expected DDoS Routing Engine violations are not observed on the MPC10E. PR1579319
- On the MX2020 routers, the `ISSU RECONNECT TIMEOUT` error message is observed on the MPC6E line crads due to which the dark window size is more than expected. PR1580658
- On MX960 routers, the R0 overlay ping fails with an error message containing the tunnel source and destination address and information about the VNI. PR1580918
- On the MX480 routers, the STP topology changes after ISSU with VSTP configuration. PR1581080
- The interface might not be added to BD after the VLAN change. PR1504374
- On the MPC10 and MPC11 line cards, the following error messages occur:

  ```
  ztchip_mqss_wanio_stream_out_disable: Waiting for available credits value to become initial
  value for WO connection failed - status 29, wan_port_group 0, conn_num 10
  ```

  PR1497089
- The MX150 Series routers might go into the `Database` mode after the software upgrade or downgrade. PR1510892
- The `show configuration` command does not display the actual version information. PR1517231
- Receipt of the specific packets might lead to Denial of Service in the MQTT Server. PR1522265
- Packets might drop with all commit events with the 1G speed configured interface. PR1524614
- Memory leaks while querying the aggregated Ethernet interface statistics. PR1528605
- The dcpfe process might crash and cause FPC to restart due to the traffic burst. PR1534340
- The MX240 routers with NG-RE reports mixed primary and backup Routing Engine types alarm.

  PR1536184
- A specific BGP VPNv6 flowspec message causes the routing protocol daemon (rpd) process to crash with a core. PR1537085
- Any modification made in the middle of the existing firewall filter might lead to all host-bound traffic being discarded. PR1544502
- The ancpd process generates core cile while hitting the maximum-discovery-table-entries limit. PR1544746

- The `jnxDomAlarmSet` and `jnxDomAlarmClear` traps do not get generated at 15 minutes intervals after a link on the transceivers support DOM becomes up or down. PR1545514
- Receipt of the specific DHCPv6 packet might cause the jdhcpd process to crash and restart. PR1546166
- The OSPFv3 session might flap and OSPFv3 hellos might drop in the host path. PR1547032
- sFlow requests for license upon committing the configuration. PR1550140
- The IRB interface might not work after chassisd and l2ald reboots in an EVPN scenario. PR1551631
- The `action-shutdown` command of the storm control does not work for the ARP broadcast packets. PR1552815
- On the MX204 and MX10003 routers, the chassisd process might crash with repeated configuration commits. PR1555271
- Client authentication fails after performing GRES. PR1563431
- Need to improve the handling deletion of static demux interface with active subscribers. PR1570739
- The following commit failure error message appears:

```
Modified IFD "ae0" is in use by targeted BBE subscriber, commit denied - mtu config changed
(1522), (1514)
```

PR1577007

Infrastructure

- Invalid statistics value might be observed when multiple mib2d/cosd requests for the same logical interface arrives within one second. PR1541579

Interfaces and Chassis

- The configuration might not be applied after deleting all existing logical interfaces and adding a new logical interface for a physical interface (IFD) in a single commit. PR1534787
- The following errors are generated during GRES: `VRRPMAN_PATRICIA_GROUP_ADD_FAIL: vrrp_ifcm_send_bulk: Failed to add group to patricia tree key` and `VRRPMAN_ENTRY_KEY_PRESENT: vrrp_ifcm_send_bulk: Already an entry present with the key`. PR1575689
- Inline Y.1731 SLM or DM does not work in the enhanced-cfm-mode for the EVPN up MEP scenario. PR1537381

- The following error message might be seen after commit for configuration under interface hierarchy:

```
should have at least one member link on a different fpc.
```

PR1539719

- The following commit error is observed while trying to delete unit 1 logical systems interfaces: ae2.1:

```
Only unit 0 is valid for this encapsulation.
```

PR1547853

- The startup-silent-period command might not work in Junos OS Release 20.3R1 or later. PR1548464
- The VCP port is marked as administratively down on the wrong MX-VC member. PR1552588
- The dcd process might leak memory on pushing the configuration to the ephemeral database. PR1553148
- On the MX960 routers, sessions are flapped after applying the action profile on the router. PR1561044
- The input errors counter on the monitor interface CLI does not work. PR1561065
- MAC address entry issue might be seen after the MC-LAG interface fails or falls back. PR1562535
- Traffic loss might be seen while verifying VRRP State Machine functionality. PR1564551
- The traceroute Local Privilege Escalation vulnerabilities in SUID binaries appears. PR1529209
- The ppmd process might crash when you configure VRRP. PR1561281
- The MC-AE interfaces might go down if you configure same VRRP group-id on the multiple IRB units. PR1575779

Juniper Extension Toolkit (JET)

- TCP connection might not be established while creating the default gRPC channel with `fw_channel` name. PR1559064

Layer 2 Ethernet Services

- The copying of files to the RCB over WAN ports is slow. PR1496895

- Receipt of the malformed DHCPv6 packets causes the jdhcpd process to crash and restart. PR1564434
- The `show dhcp relay statistics` command displays `DHCPLEASEUNASSIGNED` instead of `DHCPLEASEUNASSINGED`, which is a spelling error. PR1512239
- DHCP packet might drop when DHCP relay is configured on the leaf device. PR1554992
- jnxJdhcpLocalServerMacAddress (.1.3.6.1.4.1.2636.3.61.61.1.4.3) returns incorrect format of MAC address. PR1565540
- The Option 82 information are incorrectly cleared by the DHCP Relay Agent. PR1568344
- Receipt of a crafted DHCP packet causes the jdhcpd DHCP service to generate core files. PR1534814

MPLS

- Traffic loss might be observed due to rpd crash in the MPLS scenario. PR1528460
- MPLS LSP on transit has double entries. PR1533161
- The rpd process might crash when the LDP route with indirect next hop is deleted on the aggregated Ethernet interface. PR1538124
- Committing might trigger externally provisioned LSP MBB mechanism. PR1546824
- A new LSP might not be up even if the bypass LSP is up and `setup-protection` is configured. PR1555774

Multicast

- FPC might crash in a multicast scenario. PR1569957

Network Management and Monitoring

- Commit error while deleting the routing instance when SNMP trap-group also has the same routing instance referred. PR1555563
- The trace-relay process generates core files. PR1556040
- After the l2cpd service is restarted, the context of registration from l2cpd to snmpd was failing due to incorrect reinitialization. Because of this, if an NMS polls the dot1dStp objects by prefixing the context might fail. As a workaround, restart snmpd or reconfigure the protocols hierarchy. PR1561736

Platform and Infrastructure

- PE-CE OAM CFM might have issues in the aggregated Ethernet interface. PR1501656
- The following major error might cause Packet Forwarding Engine(s) to disable: `XQ_CMERROR_SCHED_L3_PERR_ERR` PR1538960
- An internal timer on the backup Routing Engine might cause an ARP storm upon GRES switchover on the new primary Routing Engine. PR1547583
- The state of the flow detection configuration might not be displayed properly if DDoS-SCFD is configured globally. PR1519887
- The following error message is observed when the alarms resets after interface:

  ```
  7836 ifl 567 chan_index 8 NOENT & jnh_ifl_topo_handler_pfe(13015): ifl=567 err=1 updating
  channel table nexthop.
  ```

  PR1525824
- The VXLAN encapsulation over IPv6 underlay might not work. PR1532144
- The PPE error messages or traps might be observed in the Layer 2 flooding scenarios. PR1533767
- The fpc process might crash when the next-hop memory of ASIC is exhausted in an EVPN-MPLS scenario. PR1533857
- The ISSU might fail on platforms running Junos OS with LU chip-based line cards. PR1535745
- Subscribers do not come up with VPLS on ps interface. PR1536043
- Packet loss might be observed when the RFC2544 egress reflector session is configured on the nonzero Packet Forwarding Ethernet interface. PR1538417
- The vmxt_lnx process generates a core file at `l2_metro_bd_host_inject_del bd_platform_delete bd_handle_msg`. PR1538516
- The rmopd process might leak memory if the TWAMP client is configured. PR1541808
- FPC might crash when the underlying Layer 2 interface for ARP over IRB interface is changed from the physical interface to the LSI interface. PR1542211
- ARP expired timer on the backup Routing Engine is not the same as on the primary Routing Engine if aging-timer is configured. PR1544398
- The kernel might crash if GRES is performed in either a new iteration or after swapping the Routing Engine and restoring the HA configuration. PR1549656

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. PR1552603
- Traffic is not forwarded over IRB to l2circuit on lt interfaces. PR1554908
- IPv4 EXP rewrite might not work properly when inet IPv6-VPN is enabled. PR1559018
- DHCPv4 request packets might be wrongly dropped during DDoS attacks. PR1562474
- The `enforce-strict-scale-limit-license` configuration enforces subscriber license incorrectly and the following error message is observed:

  ```
  / PADS:"AC-System-Error - No resources"
  ```

  PR1563975
- The BUM frame might be duplicated on an aggregate device if the extended-port on the Satellite device is an aggregated Ethernet interface. PR1560788
- The `IIF-LIST APP DWORD` leak is observed during the period of churn for the NGMVPN-MoFRR routes with sender-based-rpf enabled. PR1548806
- DDoS LACP violation occurs upon receipt of specific layer 2 frames in an EVPN-VXLAN scenarios. PR1512033
- During Routing Engine switchover the new primary Routing Engine might suddenly crash. PR1527246
- In a rare occurrence, the Routing Engine kernel might crash while handling the TCP sessions if you enable GRES or NSR. PR1546615
- The toe_lu_stats_ucode process generates the core file at `jbeta_fcv_alloc_fcv_idx_global jbeta_sfilter_fcv_cb bwy_dfw_sfilter_fcv_cb`. PR1569328
- On the MX480 routers with Trio line cards hosting subscribers, when memory allocation in the counter segment fails the error logs might be observed. PR1570631
- On the MX480 routers, FPC reports following error log message:

  ```
  cassxr_err_addr(8593): Uninitialized Read Error @ EDMEM[0x7cb601b0]
  ```

  PR1573920

Routing Policy and Firewall Filters

- For setting the IPv6 router ID, the `routing-options` statement is added. PR1523283
- The RPD process generates core file at `task_block_alloc_jemalloc isis_spring_stats_jobinfo_alloc isis_spring_stats_show_traffic_stats`. PR1579830
- The policy configuration might be mismatched between the rpd and mgd processes when deactivate policy-options prefix-list is involved in the configuration sequence. PR1523891
- The generated route goes into the **Hidden** state when the `protect core` statement is enabled. PR1562867
- Global variable policy_db_type do not set the correct value on failure. PR1561931

Routing Protocols

- The BFD session might get stuck in the **Init** or **Down** state after the BFD session flaps. PR1474521
- With BGP rib-sharding enabled, RPD memory might exhaust. PR1546347
- Traffic might be lost during mirror data transmit from the primary ppmd/bfdd. PR1570228
- VRF table does not get refreshed after changing to `maximum-prefixes` in the VRF. PR1564964
- Traffic loss might occur during VRF route resolution over indirect nexthop. PR1525363
- The rpd might crash with BGP RPKI enabled in a race condition. PR1487486
- The virtual-router option is not supported under a routing instance in a lean rpd image. PR1494029
- Some PIM join or prune packets might not be processed in the first attempt in the scale scenario where the PIM routers establish neighborship and immediately join the multicast group. PR1500125
- Traffic might be silently discarded when the clear bgp neighbor all command is executed on a router and also on the corresponding route reflector in succession. PR1514966
- The BGP session with VRRP virtual address might not come up after a flap. PR1523075
- The VRF label is not assigned at ASBR when the inter-AS is implemented. PR1523896
- The rpd process generates a core file at `is_srv6_delete_locator_end_sid_data isis_srv6_end_sid_local_data_delete isis_srv6_locator_config_check`. PR1531830
- Transit labels for Layer 3 VPN routes are pushed momentarily to the MPLS.0 table. PR1532414
- Configuring then next hop and then reject on a route policy for the same route might cause the rpd process to crash. PR1538491

- After the peer is moved out of the protection group, the path protection is not removed from the PE device. The multipath route is still present. PR1538956
- The rpd process generates a core file at `gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash`. PR1541768
- Traffic loss might be seen in next-hop-based dynamic tunnels of the Layer 3 VPN scenario after changing the dynamic-tunnel preference. PR1542123
- Continuous rpd crash might be observed if a static group is added to the PIM protocol. PR1542573
- The metric of prefixes in intra-area-prefix LSA might be changed to 65535 when the metric of one of the OSPFv3 P2P interfaces is set to 65535. PR1543147
- IS-IS does not call ted_add_halflink for P2P IPv6-only links for traffic engineering topology. PR1548506
- Telemetry key value for transport or remote-address field for link-local IPv6 peer is incorrect, and logical interface is absent. PR1548754
- The BGP session neighbor shutdown configuration does not affect the non-established peer. PR1554569
- The changes are not effective when the values are set under static default hierarchy. PR1555187
- The BGP session might not come up if `extended-nexthop` is enabled by default on the other vendor remote peer. PR1555288
- Sending multicast traffic to downstream receiver on Virtual Chassis platforms might fail. PR1555518
- Six PE device prefixes might not be removed from the RIB upon reception of withdrawal from a BGP neighbor when RIB sharding is enabled. PR1556271
- Multipath information still shown for BGP route even after disabling interface for one path. PR1557604
- Extra node-spring-algorithm-type is displayed under the `show route table lsdist.0 te-node-iso <> extensive` command. PR1560003
- VPN routes learned from core were not advertised to the CE devices when BGP sharding is configured. PR1560661
- All Layer 3 VPN route ages reset when a VRF is added or deleted. PR1560827
- Duplicate LSP next hop is shown on inet.0, inet.3, and mpls.0 route table when OSPF traffic-engineering shortcuts and MPLS `bgp-igp-both-ribs` are enabled. PR1561207

- Wrong SPF calculation might be observed for OSPF with ldp-synchronization hold-time configured after interface flaps. PR1561414
- BGP routes might be stuck in routing table in the **Accepted DeletePending** state when the BGP peering session goes down. PR1562090
- The rpd might crash on the backup Routing Engine after rpd restart is triggered on the primary Routing Engine. PR1563350
- SNMP MIB OSPFv3NbrState returns a drifted value. PR1571473
- The rpd process crashes when you configure a fresh router with IS-IS and RIB-group to leak the inet3 routes from the no-forwarding to primary instance in a single commit. PR1534486
- The rpd process might crash when a BGP session re-establishes or flaps. PR1567182
- There might be 10 seconds delay to upload the LSP on the point-to-point interface if the rpd process restarts on its direct neighbor. 1571395
- The ppmd process might crash when you enable the MD5 authentication on OSPF with BFD flapping. PR1576893
- The rpd process generates the core file at `thread_next_node jnx_bgp_tunnel_encaps_attr_tunnel_count jnx_bgp_tunnel_encaps_attr_set_tunnel`. PR1579818

Services Applications

- L2TP subscribers might fail to establish a session on the MX Series device if the CPE is a virtual host. PR1527343
- The following error message is observed:

  ```
  SPD_CONN_OPEN_FAILURE: spd_pre_fetch_query: unable to open connection to si-1/0/0.
  ```

  PR1550035
- Executing CLI command repetitively might cause the system to run out of disk space. PR1537772

User Interface and Configuration

- The `verbose` command unexpectedly becomes hidden after Junos OS Release 16.1 for `set system export-format json`. PR1547693
- The `request system software validate on host` command does not validate the correct configuration file. PR1553577

- The configuration under groups stanza is not inherited properly. PR1529989
- Removing the flash component from Monitor > Interfaces and DHCP pages, removes the other flash pages. PR1553176
- The firewall filter for both IPv4 and IPv6 might not work when it is applied through apply-groups. PR1534858
- The JNH memory might leak on the Trio-based line cards. PR1542882

VPNs

- The PIM (S,G) join state might stay forever when there are no MC receivers and the source is inactive. PR1536903
- MVPN multicast route entry might not be properly updated with the actual downstream interfaces list. PR1546739
- Selective multicast tunnel (S-PMSI) fails to come up due to incorrect community. PR1537636
- Type 7 messages might not be sent from the egress PE devices resulting in the Type 3 or Type 5 messages not created for some S, Gs in the source PE devices. PR1567584

Documentation Updates

There are no errata and changes in Junos OS Release 21.1R3 for the MX Series documentation.

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 17.4R1 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 21.1R3

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

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

The installation process rebuilds the file system and completely reinstalls Junos OS. Configuration information from the previous software installation is retained, but the contents of log files might be erased. Stored files on the routing platform, such as configuration templates and shell scripts (the only exceptions are the juniper.conf and ssh files might be removed. To

preserve the stored files, copy them to another system before upgrading or downgrading the routing platform. For more information, see the Installation and Upgrade Guide.

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

Procedure to Upgrade to FreeBSD 11.x-Based Junos OS

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

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

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

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

2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the release number (the number of the software version that you want to download) from the Release drop-down list to the right of the Download Software page.
4. Select the Software tab.
5. In the Install Package section of the Software tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by a Juniper Networks representative.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.
9. Copy the software to the routing platform or to your internal software distribution site.
10. Install the new jinstall package on the routing platform.

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

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

- For 32-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-x86-32-21.1R3.9-signed.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-21.1R3.9-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-x86-32-21.1R3.x-limited.tgz
```

- For 64-bit Routing Engine version:

```
user@host> request system software add no-validate reboot source/junos-install-mx-x86-64-21.1R3.9-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.
- Starting in Junos OS Release 21.1R3, in order to install a VM host image based on Wind River Linux 9, you must upgrade the i40e NVM firmware on the following MX Series routers:
  - MX240, MX480, MX960, MX2010, MX2020, MX2008, MX10016, and MX10008

  [See https://kb.juniper.net/TSB17603.]

NOTE: After you install a Junos OS Release 21.1R3 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-21.1R3.9-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-21.1R3.9-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 21.1R3 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 and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

  Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 13: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

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 21.1R3

To downgrade from Release 21.1R3 to another supported release, follow the procedure for upgrading, but replace the 21.1R3 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.

Junos OS Release Notes for NFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for the NFX Series Network Services Platforms. 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for the NFX Series.

What's New in 21.1R3

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features or enhancements to existing features for NFX Series devices in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for the NFX Series.

Application Identification (AppID)

- **Application signature package enhancements (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 21.1R1, we've enhanced the application signature package by grouping all newly added signatures under the junos:all-new-apps group. When you download the application signature package on your device, the predefined application group is downloaded. You can use this application group in the security policy configuration.

We've also introduced a list of application tags, based on attributes, in the application signature package. You can group similar applications based on these predefined tags. By doing so, you can consistently reuse the application groups when you define security policies.

[See Predefined Application Signatures for Application Identification.]

- **Enhancements to packet capture of unknown applications (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 21.1R1, your security device stores the packet capture of unknown applications' details per session. As a result of this change, the packet capture (.pcap) file now includes the session ID in the filename. We now store the file in **destination-IP-address.destination-port.protocol.session-id.pcap** format in the **/var/log/pcap** location. (Previously, the packet capture file was saved in **destination-IP-address. destination-port.protocol.pcap** format.)

  In addition, we've enhanced packet capture of unknown application functionality to capture unknown Server Name Indication (SNI) details.

  [See Packet Capture of Unknown Application Traffic Overview.]

- **Application signature enhancements (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 21.1R1, we've introduced the following enhancements to application signatures:
  - Support for FTP data context propagation
  - Skipping of deep packet inspection (DPI) for the sessions offloaded by advanced policy-based routing (APBR) on application system cache (ASC) hit (when only APBR service is enabled).
  - Forceful installation of the application signature pack over the same version of signature pack.
  - Display (in the CLI command output) of the application signature pack release date.
  - Display (in the CLI command output) of the list of deprecated application signatures available in the installed signature pack.

  [See Predefined Application Signatures for Application Identification.]

Architecture

- **Custom mode (NFX250 NextGen and NFX350 devices)**—Starting in Junos OS Release 21.1R1, you can define and specify a custom-mode template for NFX250 NextGen and NFX350 devices. The custom mode provides an option to allocate resources to Layer 3 data plane and Network Functions Virtualization (NFV) backplane.
  [See NFX350 Overview and NFX250 NextGen Overview.]

Flow-Based and Packet-Based Processing

- **Support for PowerMode IPsec (PMI) solution (SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800 with SPC3 cards, vSRX, and vSRX3.0) and GRE acceleration solution (SRX Series and NFX Series)**—Starting in Junos OS Release 21.1R1, we support the PMI and GRE acceleration solutions to improve the software-defined WAN (SD-WAN) performance.

  Table 14: Solutions and Details

Solution	How to Enable?
PMI	Include the `power-mode-ipsec` and `gre-performance-acceleration` statements at the `[edit security flow]` hierarchy level. **NOTE**: PMI supports both IPsec and GRE. In this case, traffic flows through the PMI data path.
GRE acceleration	Include the `gre-performance-acceleration` statement at the `[edit security flow]` hierarchy level. **NOTE**: By default, `gre-performance-acceleration` is turned off. In this case, traffic flows through the GRE acceleration data path.

  [See gre-performance-acceleration (Security Flow), flow (Security Flow), and show security flow status.]

Intrusion Detection and Prevention

- **Support for Perl-compatible regular expression (PCRE) version 8.40 (SRX Series and NFX Series)**—Starting in Junos OS Release 21.1R1, we've upgraded the codebase of intrusion detection and prevention (IDP) from PCRE version 5.40 to PCRE version 8.40. As PCRE version 8.40 supports new regex constructs, this upgrade enhances the capability of Junos OS IDP attack signatures to match regular expressions. With this upgrade, we've also addressed security vulnerabilities in the Junos OS PCRE codebase.

  [See pattern-pcre (Security IDP).]

- **Support for Snort IPS signatures (SRX Series and NFX Series)**—Starting in Junos OS Release 21.1R1, Juniper Networks IDP supports Snort IPS signatures. IDP secures your network by using signatures that help to detect attacks. Snort is an open-source intrusion prevention system (IPS). You can convert the Snort IPS rules into Juniper IDP custom attack signatures using the Juniper Integration of Snort Tool (JIST). These rules help detect malicious attacks.

- JIST is included in Junos OS by default. The tool supports Snort version 2 and version 3 rules.
- JIST converts the Snort rules with snort-ids into equivalent custom attack signatures on Junos OS with respective snort-ids as the custom attack names.
- When you run the `request` command with Snort IPS rules, JIST generates `set` commands equivalent to the Snort IPS rules. Use the `request security idp jist-conversion` command to generate the `set` commands as CLI output. To load the `set` commands, use the `load set terminal` statement or copy and paste the commands in the configuration mode, and then commit. You can then configure the existing IDP policy with the converted custom attack signatures.
- All the Snort IPS rule files that didn't get converted are written to **/tmp/jist-failed.rules**. The error log files generated during the conversion are written to **/tmp/jist-error.log**.
- To view the jist-package version, use the `show security idp jist-package-version` command.

[See Understanding Snort IPS Signatures, request security idp jist-conversion , and show security idp jist-package-version .]

Platform and Infrastructure

- **Transfer files from USB (NFX150, NFX250 NextGen, and NFX350 devices)**—Starting in Junos OS Release 21.1R1, you can transfer files from USB to NFX devices by enabling the USB pass-through feature. To enable this feature, use the `set system services usb-pass-through` command. Built-in LTE functionality does not work after you enable the USB pass-through feature.

  [See Supporting File Transfer from USB on NFX150 Devices, Supporting File Transfer from USB on NFX250 NextGen Devices, and Supporting File Transfer from USB on NFX350 Devices.]

- **Virtual port peering (NFX250 NextGen and NFX350 devices)**—Starting in Junos OS Release 21.1R1, you can configure the virtual port peering (VPP) feature to map a physical port and an interface to a virtualized network function (VNF), so that if the physical interface becomes inactive, the corresponding virtual interface also becomes inactive and the status of the physical interface is communicated to the virtual interface.

  The VPP feature is supported only on the Network Functions Virtualization (NFV) backplane.

  [See Configuring VNFs on NFX350 Devices and Configuring VNFs on NFX250 NextGen Devices.]

What's Changed

IN THIS SECTION

Learn about what changed in Junos OS main and maintenance releases for NFX Series devices.

What's Changed in Release 21.1R3

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in this release for NFX Series devices.

What's Changed in Release 21.1R2

There are no changes in the behavior of Junos OS features or in the syntax of Junos OS statements and commands in this release for NFX Series devices.

What's Changed in Release 21.1R1

IN THIS SECTION

Network Management and Monitoring

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

[See ssh (NETCONF).]

Known Limitations

There are no known limitations in hardware or software in Junos OS Release 21.1R3 for NFX Series devices.

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

Open Issues

IN THIS SECTION

Learn about open issues in Junos OS Release 21.1R3 for NFX Series devices.

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

High Availability

- On an NFX350 chassis cluster, when FPC0 (when node0 is primary) or FPC7 (when node1 is primary) is restarted by either using the `request chassis fpc slot slot restart node local` command or because of dcpfe core files on the primary, it restarts FPC1 or FPC8. This might break the pre-existing TCP sessions and fail to restart the TCP sessions. The TCP sessions might require a manual restart. PR1557607

Platform and Infrastructure

- On NFX150 devices, the following error message appears during FTP: `ftpd[14105]: bl_init: connect failed for `/var/run/blacklistd.sock' (No such file or directory)`. PR1315605

Resolved Issues

IN THIS SECTION

Learn about the issues fixed in these releases 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: 21.1R3

IN THIS SECTION

Interfaces

- Unable to configure destination-port on firewall filter on NFX250 NextGen devices. PR1592019
- On NFX Series devices, deletion of VNF interfaces that are mapped SR-IOV interface fails. PR1598993
- L3 data plane interfaces are not appearing when flex mode is enabled on NFX350-S3 devices. PR1599643

User Interface and Configuration

- When the available free physical memory drops below 1.5 GB, configuration commits by Junos Device Management Daemon (JDMD) might not take effect and mustd core files will be seen. This will not have any impact on the running traffic. PR1599641

Resolved Issues: 21.1R2

IN THIS SECTION

General Routing

- RPD core file is generated when the device reboots and daemon restarts. Daemon recovers and there is no service impact on routing protocol usage. PR1567043

Interfaces

- AE interface statistics are not reported on NFX250 devices. PR1581596
- LACP subsystem is not enabled in NFX250 NextGen devices. PR1581717
- On NFX Series devices, you need to adjust MTU sizes of the OVS system interfaces to maintain consistency. PR1586967
- On NFX Series devices, deletion of VNF interfaces that are mapped SR-IOV interface fails. PR1598993

Performance Modes

- You cannot enable the trust mode on an SR-IOV virtual function assigned to a VNF. PR1593037

Platform and Infrastructure

- You can transfer file from USB to hypervisor by enabling the usb-pass-through functionality. PR1535220
- On NFX150 devices, when J-Flow v5 is configured and the J-Flow v5 server is reachable through anIPsec tunnel, and the MTU size of this IPsec tunnel is configured as 1500, the J-Flow packets are not generated on NFX Series devices. PR1539964
- Zeroise is successful in NFX350 devices. However, an error message is reported during zeroize operation. PR1565077

Resolved Issues: 21.1R1

IN THIS SECTION

High Availability

- On NFX150 devices, upgrade from Junos OS Release 19.4 to Junos OS Release 20.2 fails and the `/usr/sbin/boot_mgmt_fsm: line 40: echo: write error: No space left on device issue` message is displayed. PR1532334

Interfaces

- On NFX250 devices, a VNF interface is not brought down when the VNF interface is mapped to an already link down or disabled peer physical interface. PR1555193
- Analyzer on OVS fails to mirror packets after a system reboot on a DPDK-enabled device. PR1480290
- On NFX Series devices, the following error message for interfaces might be seen: `FAILED(-1) read of SFP eeprom.` PR1529939

Performance Modes

- A message is provided in syslog if reboot is required for the mode modification to take effect in custom mode. PR1555465

Platform and Infrastructure

- On NFX150, NFX250 NextGen, and NFX350 devices, the `EmulatorPin CPUSet` option does not get configured, which might result in vCPU running on a higher level up to 100%. PR1540564
- On NFX350 devices and the SRX5000 line of devices with SPC3 card, the DPD Gateway failover feature is not supported. PR1564715
- The l2cpd core files might be seen on reboot. PR1561235
- The DSL SFP firmware cannot finish upgrade successfully through vmhost reboot. PR1547540

Documentation Updates

There are no corrections or changes in Junos OS Release 21.1R3 documentation for the NFX Series documentation.

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.

NOTE: For information about NFX product compatibility, see NFX Product Compatibility.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

  Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 15: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Basic Procedure for Upgrading to Release 21.1

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 Software Installation and Upgrade Guide.

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

To download and install Junos OS Release 21.1R2

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

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

2. Select the name of the Junos OS platform for the software that you want to download.
3. Select the **Software** tab.
4. Select the release number (the number of the software version that you want to download) from the Version drop-down list to the right of the Download Software page.
5. In the Install Package section of the Software tab, select the software package for the release.
6. Log in to the Juniper Networks authentication system using the username (generally your e-mail address) and password supplied by Juniper Networks representatives.
7. Review and accept the End User License Agreement.
8. Download the software to a local host.

9. Copy the software to the device or to your internal software distribution site.

10. Install the new package on the device.

Junos OS Release Notes for PTX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for the PTX Series.

What's New in 21.1R3

There are no new features or enhancements to existing features in Junos OS Releases 21.1R3, 21.1R2, or 21.1R1 for PTX Series.

What's New in 21.1R2

There are no new features or enhancements to existing features in this release for the PTX Series.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for the PTX Series.

High Availability

- **Support for VRRP (PTX1000, PTX10002, PTX10008, and PTX10016)**—Starting in Junos OS Release 21.1R1, PTX1000, PTX10002, PTX10008, and PTX10016 routers support VRRP. However, these routers do not support the following VRRP features:
  - VRRP on IRB
  - Dual tagging
  - GRES
  - VRRP on logical tunnel (LT) interfaces
  - Layer 2 VRRP

[See Understanding VRRP.]

MPLS

- **Nonstop active routing (NSR) support for controller-initiated RSVP label-switched paths (LSPs) (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 21.1R1, we support NSR for controller-initiated RSVP-based point-to-point (P2P) and point-to-multipoint (P2MP) LSPs. The primary Routing Engine synchronizes all RSVP LSPs initiated by Path Computation Elements (PCEs), including multicast flow specifications for any PCE-initiated P2MP LSPs, with the backup Routing Engine. This ensures zero traffic loss for the traffic carried over PCE-initiated RSVP LSPs during Routing Engine switchovers. This feature is enabled when NSR is configured.

  [See PCEP Configuration.]

- **BGP Classful Transport planes (BGP-CT) to facilitate service mapping over colored tunnels (ACX Series, PTX Series, MX Series)**—Starting in Junos OS Release 21.1R1, you can classify colored transport tunnels (RSVP, IS-IS flexible algorithm) in your network into transport classes and map service routes over an intended transport class. You can also extend the transport tunnels to span across multiple domains (ASs or IGP areas) by using the new BGP transport address family called BGP Classful Transport (BGP CT).

  This feature lays the foundation for network slicing and allows the different domains to interoperate irrespective of the transport signaling protocols used in each domain.

  [See BGP Classful Transport Planes Overview.]

Multicast

- **Controller-based BGP multicast signaling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 21.1R1, we've introduced controller-based BGP multicast signaling that can be used instead of hop-by-hop signaling to program multicast forwarding states on routers. An external controller that is aware of the topology and network events within that topology calculates the optimum multicast trees between the source and receivers. The external controller then uses BGP signaling to send a new type of BGP network layer reachability information (NLRI) with modified attributes to convey the multicast state information to all the routers on the multicast trees.

  You can use this feature instead of multicast routing protocols, such as Protocol Independent Multicast (PIM) or multipoint LDP (MLDP). You can enable this feature using `bgpmcast` configuration option at the `[edit protocols]` hierarchy.

Network Management and Monitoring

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448,**

EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

[See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

- **sFlow support for IP-IP traffic (PTX1000, PTX10008, and QFX10002)**—Starting in Junos OS Release 21.1R1, you can use sFlow technology to sample IP over IP (IP-IP) traffic on a physical port. sFlow sampling is supported for IP-IP tunnels that have an IPv4 outer header that carry IPv4 or IPv6 traffic. You can use sFlow monitoring technology to randomly sample network packets from IP-IP tunnels and to send the samples to a destination collector for monitoring. Devices that act as an IP-IP tunnel entry point, transit device, or tunnel endpoint support sFlow sampling.

  [See Overview of sFlow Technology and Configuring IP Tunnel Interfaces.]

Routing Protocols

- **Support for configuring multiple independent IGP instances of IS-IS (ACX Series, MX Series, and PTX Series)**—Starting in Junos OS Release 21.1R1, you can configure and run multiple independent IGP instances of IS-IS simultaneously on a router.

  > **NOTE**: Junos OS does not support configuring the same logical interface in multiple IGP instances of IS-IS.

  [See How to Configure Multiple Independent IGP Instances of IS-IS.]

- **Support for IP forward backup path for BGP-LS peer SIDs (PTX Series)**—Starting in Junos OS Release 21.1R1, you can configure an IP forward backup path that provides protection at the local node or the point of local repair for egress peer engineering. When the primary segment goes down, the packet is forwarded to the configured IP backup path. This IP forward backup path has local node significance only. BGP does not send the IP forward backup path information to the controller in its periodic BGP Link State (BGP-LS) updates. If you have configured both segment protection and IP forward backup path, then backup segment protection takes precedence over the IP forward backup path protection.

To configure IP forward backup path for BGP-LS peer segments, include the `egress-te-backup-ip-forward` option at the `[edit bgp egress-te-segment-set]`, `[edit bgp group group-name egress-te-node-segment]`, and `[edit bgp group group-name egress-te-segment adj]` hierarchy levels.

[See egress-te-set-segment, ,egress-te-node-segment, and egress-te-adj-segment.]

- **Support for BGP Auto-discovered Neighbor (MX Series, PTX1000, PTX10008, QFX5120-32C, QFX5200, QFX5210, and QFX10008)**—Starting in Junos OS Release 21.1R1, we support BGP auto-discovered neighbors using IPv6 Neighbor Discovery Protocol (ND). With this feature, you can enable BGP to create peer neighbor sessions using link-local IPv6 addresses of directly connected neighbor devices. You need not specify remote or local neighbor IP addresses.

  To enable peering for a given interface or set of interfaces without specifying the local or remote neighbor addresses, configure the `peer-auto-discovery` statement at the `[edit fabric protocols bgp group <name> dynamic-neighbor <name>]` hierarchy level.

  [See BGP Auto-Discovered Neighbors, and peer-auto-discovery.]

Segment Routing

- **Support for flexible algorithm in OSPFv2 for segment routing traffic engineering (ACX5448, ACX710, MX204, MX104, MX480, MX960, MX10003, MX2020, and PTX10001)**—Starting in Junos OS Release 21.1R1, you can thin-slice a network by defining flexible algorithms that compute paths using different parameters and link constraints based on your requirements. For example, you can define a flexible algorithm that computes a path to minimize IGP metric and define another flexible algorithm to compute a path based on traffic engineering metric to divide the network into separate planes. This feature allows networks without a controller to configure traffic engineering and utilize segment routing capability of a device.

  To define a flexible algorithm, include the `flex-algorithm` statement at the `[edit routing-options]` hierarchy level.

  To configure a device to participate in a flexible algorithm, include the `flex-algorithm` statement at the `[edit protocols ospf source-packet-routing]` hierarchy level.

  [See How to Configure Flexible Algorithms in OSPF for Segment Routing Traffic Engineering.]

- **Support for strict SPF and IGP shortcut (ACX710, MX960, MX10008, MX2020, PTX5000, and PTX1000)**—Starting in Junos OS Release 21.1R1, you can configure segment routing algorithm 1 (strict SPF) and advertise its SIDs in IS-IS link-state PDU (LSPDU) and use these SIDs to create SR-TE tunnels to forward the traffic by using the shortest IGP path to reach the tunnel endpoint while avoiding loops. You can also specify a set of prefixes in the import policy, based on which the tunnel can redirect the traffic to a certain destination. You can use algorithm 1 (strict SPF) along with algorithm 0 (default SPF) by default when Source Packet Routing in Networking (SPRING) is enabled.

  [See How to Enable Strict SPF SIDs and IGP Shortcut, prefix-segment, and source-packet-routing.]

Services Applications

- **TWAMP Light IPv4 support (MX Series, PTX Series)**—Starting in Junos OS Release 21.1R1, we support the Two-Way Active Measurement Protocol (TWAMP) Light, as defined in Appendix I of RFC 5357. TWAMP Light is a stateless version of TWAMP, where test parameters are predefined instead of negotiated. All test packets received by the server on a test port are reflected back and forgotten right away.

  [See twamp.]

- **Support for inline active flow monitoring (PTX10008 and PTX10016)**—Starting in Junos OS Release 21.1R1, we support inline active flow monitoring for the PTX10K-LC1105 line card. Inline active flow monitoring supports version 9 and IPFIX flow collection templates. Both the IPFIX and the version 9 templates are supported for IPv4, IPv6, and MPLS, and use UDP as the transport protocol.

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

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for the PTX Series.

What's Changed in Release 21.1R3

IN THIS SECTION

Junos XML API and Scripting | 259

EVPN

- **Output for show Ethernet switching flood extensive**—The output for the `show ethernet-switching flood extensive` command now displays the correct next-hop type for Virtual Ethernet and WAN mesh group in an EVPN-VXLAN network as unilist. Previously, the output for the `show ethernet-switching flood extensive` command would misidentify the next-hop type as composite.

General Routing

- **No support for PKI operational mode commands on the Junos Limited version (MX Series routers, PTX Series routers, and SRX Series devices**—We do not support `request` , `show` , and `clear` PKI-related operational commands on the limited encryption Junos image ("Junos Limited"). If you try to execute PKI operational commands on a limited encryption Junos image, then an appropriate error message is displayed. The `pkid` process does not run on Junos Limited version image. Hence, the limited version does not support any PKI-related operation.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Prior to this PR, the source and destination address pair had to be unique among all the FTI tunnels regardless of the tunnel encapsulation type.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See request system scripts refresh-from and cert-file (Scripts).]

What's Changed in Release 21.1R2

There are no changes in behavior and syntax for PTX Series in Junos OS Release 21.1R2.

What's Changed in Release 21.1R1

IN THIS SECTION

General Routing

- **Change in severity of fabric output CRC errors (PTX5000)**—We've reduced the severity of fabric output CRC errors from fatal to minor. With this change, the fabric output CRC errors (CMERROR_TQ_FO_CRC) no longer cause the Packet Forwarding Engines to be disabled.
- **Secure boot disabled alarm is raised (PTX10008)**—The `Secure boot disabled` alarm is raised when the system boots with secure boot disabled in bios.

Interfaces and Chassis

- **Warning message when taking an FPC offline**—PTX10003-80C and PTX10003-160C devices do not support the `request chassis fpc slot` *`slot-number`* `online` command. The only way to bring up an FPC (MPC) that is offline is by rebooting the chassis. So, when you take an FPC offline by using the `request chassis fpc slot` *`slot-number`* `offline` command, the screen displays the following message: `Warning : FPC <slot> cannot be made online using a CLI command. You need to perform router reboot using "request system reboot" to online the FPC <slot>. Do you wish to continue ? [yes,no] (no)`.

  [See request chassis fpc.]

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX event scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Network Management and Monitoring

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can

configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the `[edit system services netconf hello-message yang-module-capabilities]` hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the `[edit system services netconf netconf-monitoring netconf-state-schemas]` hierarchy level.

[See hello-message and netconf-monitoring.]

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the `verbose` statement at the `[edit system export-format json]` hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from `verbose` to `ietf` starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `[edit system export-format json]` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

  [See export-format.]

Known Limitations

IN THIS SECTION

Learn about known limitations in Junos OS Release 21.1R3 for PTX Series.

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

General Routing

- Duplicate IPv6 explicit-null MPLS labels might be present in the MPLS packet during the BGPoLDPoRSVP scenario. PR1556328
- On PTX10008 routers, end to end traffic does not flow for ethernet switching in the EP style. PR1583219

Open Issues

IN THIS SECTION

Learn about open issues Junos OS Release 21.1R3 for PTX Series.

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

General Routing

- On PTX Series routers with FPC-PTX-P1-A or FPC2-PTX-P1A, you might encounter a single event upset (SEU) event that might cause a linked-list corruption of the TQCHIP. The following syslog message gets reported:

```
Jan 9 08:16:47.295 router fpc0 TQCHIP1: Fatal error pqt_min_free_cnt is zero
Jan 9 08:16:47.295 router fpc0 CMSNG: Fatal ASIC error, chip TQ
Jan 9 08:16:47.295 router fpc0 TQ Chip::FATAL ERROR!! from PQT free count is zero
jan 9 08:16:47.380 router alarmd[2427]: Alarm set: FPC color=RED, class=CHASSIS, reason=FPC 0
Fatal Errors - TQ Chip Error code: 0x50002
```

```
Jan 9 08:16:47.380 router craftd[2051]: Fatal alarm set, FPC 0 Fatal Errors - TQ Chip Error
code: 0x50002
```

The Junos OS Chassis Management error handling detects such a condition, raises an alarm, and disables the affected Packet Forwarding Engine entity. To recover this Packet Forwarding Engine entity, restart the FPC. Contact your Juniper support representative if the issue persists even after the FPC restarts.

PR1254415

- The following log message might get generated on FPC with WINTEC mSATA SSD:

```
SMART ATA Error Log Structure error: invalid SMART checksum.
```

PR1354070

- The firewall counter for lo0 interface might not increase. PR1420560
- On PTX1000 routers, interface flaps during ZTP. PR1534614
- On PTX1000 routers, the `Extended Router Data` VRFoIPoIP headend sFlow record displays incorrect next hop in an ECMP case. PR1537190
- On PTX1000 routers, the IPoIP transit sFlow egress record displays incorrect `Extended Switch Data` when using VLAN interface. PR1537648
- On PTX1000 routers, after upgrading or downgrading the VM host platform, during booting up with the new image, the Wind River Linux (WRL) kernel might go into the `Deadlock` state due to a race condition in the Advanced Configuration and Power Interface (ACPI) Component Architecture (ACPICA) module in Linux kernel. This might cause the system to become nonresponsive in continuous crashing state. PR1544875
- On PTX10002-60C routers, when you configure an inline jFlow and set high sampling rate (more than 4000 per second), high CPU utilization might occur and this might result in relevant impacts on the traffic analysis and billing. PR1569229
- On PTX10002-60c routers, when you configure a firewall with both discard and port-mirror as actions in the same term, mirrored packet gets corrupted (have two Layer 2 headers). PR1576914
- On PTX5000 routers, IS-IS adjacency does not come up through the circuit cross-connect Layer 2 circuit. PR1590387
- In configurations where a large number of tag next-hops have ND6 (Neighbour Discovery) next-hop as underlying next-hop, upon refresh of ND6 entry because of any reason a large number of updates are sent to the Packet Forwarding Engine. This update processing causes a spike in the CPU usage, which can hamper some scheduled tasks if they coincide. PR1600318

- On the PTX Series routers with NG-RE installed, upgrading the Intel i40e-NVM firmware to version 6.01 might generate the FRUs disconnection alarms along with traffic loss. PR1529710
- When you configure a Provider Edge (PE) router with multipath, traffic loss might occur even if the link is in the `Up` state. PR1618507

Manageability

- If the `request system zeroize` command does not trigger zero-touch provisioning, you must re-initiate the ZTP as a workaround. PR1529246

Routing Protocols

- There is traffic loss when device boot up during IGP overload. PR1495435

User Interface and Configuration

- The routing policy in the normal-mode does not revert once the network-services mode gets changed to the network-services enhanced mode. PR1587174

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for PTX Series.

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

Resolved Issues: 21.1R3

IN THIS SECTION

General Routing

- On PTX10002-60C routers, high FPC CPU utilization might occur. PR1585728
- The na-grpc process might crash and the existing telemetry connections might get disconnected. PR1587956
- The l2cpd-agent might become unresponsive after starting the telemetry service. PR1592473
- On PTX1000 router, the sFlow data (inner VLAN and outer VLAN value, forwarding-class, and DSCP value) does not get exported while checking from the server flow-records at the collector for the ingress sampling. PR1598263
- CRC errors increase continuously after the interface flaps. PR1600768
- Traffic might get silently discarded due to the `RS Fatal` error on the FPC-PTX-P1-A, FPC2-PTX-P1A, FPC-SFF-PTX-P1-A, and FPC-SFF-PTX-T. PR1600935
- The l2circuit packets with destination MAC 01:00:0c:cc:cc:cd might get punted. PR1601360
- The IPv6 traffic might be impacted when an IPv6 route resolves over a dynamic tunnel. PR1602007
- On PTX10002-60C router, after upgrading the device, the configured firewall filters might be applied on the incorrect interfaces. PR1602292
- Packet loss might occur on the filter-based GRE deployments. PR1603453
- On PTX5000 routers, link might flap momentarily. PR1606008
- Memory might leak on the l2cpd process when you commit certain LLDP operations. PR1608699

- There is a one-shot timer created for LLDP(Link Layer Discovery Protocol), which might not get freed before creating the new one-shot timer because of which 160 bytes of leakage occurs every minute. This gradual memory leak in the l2cpd process might lead to the l2cpd process crash. This might impact traffic only if protocols other than LLDP (example xSTP) runs along with LLDP. PR1617151

MPLS

- The LDP replication session might not get synchronized when you enable dual-transport. PR1598174
- The VPLS connection might get down if you configure the dual-transport. PR1601854

Network Management and Monitoring

- On PTX10008 routers, syslog does not log information on the IPv4 post upgrade. PR1611504

Routing Protocols

- The rpd process might crash when you configure the BGP RPKI session record-lifetime less than hold-time. PR1585321

Resolved Issues: 21.1R2

IN THIS SECTION

EVPN

- The EVPN option under the `routing-instances <>` protocols is not present. PR1581821

Forwarding and Sampling

- The user-defined ARP policer gets applied on the aggregated Ethernet interface until the firewall process restarts. PR1528403

General Routing

- Need to improve the `request system software delete` command to add the new`archived` option to delete all the old software versions except the current and rollback. PR1566173
- Upgrading the PTX1000 devices with unified SSDs (2x32G SSD) might result in a boot loop in certain scenario. PR1571275
- The toe_gld_toe0_ucode process generates the core files at `prds_rt_ifl_ipv6_del_hndl_from_desc_list`. PR1571279
- On the PTX5000 devices, traffic loss might occur. PR1578511
- TACACS traffic might be dropped. PR1578579
- The PTX Series routers might drop traffic. PR1580211
- Configuring and deleting the FEC mode disables the auto-FEC91 on an interface that uses QSFP28-SR4 (PTX5000 or PTX3000). PR1582200
- The `show chassis clocks` command must be handled in a graceful way or with a meaningful error. PR1583715
- The following error message appears during bootup:

  ```
  Failed to get pechip handle for chip 0" and "prds_encap_sample_flood_lpbk_desc_install:
  Egress NH descriptor install OK for Flabel 7808
  ```

  PR1585594
- Traffic loss might be observed post changing the SAK keys. PR1591432
- Packet might drop on the aggregated Ethernet interface bundle with a single child member. PR1551736
- The micro BFD session might flap with the DDoS policer. PR1557782

- Device might run out of service post GRES or unified ISSU. PR1558958
- On PTX10002-60C router, another port shutdowns after shutting down one of the port. PR1568294
- Need to support LLDP Out-of-Bounds read vulnerability in l2cpd. PR1569312
- The gRPC session becomes nonresponsive in the `Closed` state. PR1571999
- On PTX10002 router, the channelized ports might drop traffic. PR1575742
- The BFD sessions might flap during traffic spikes. PR1578599
- Authentication might fail if the password contains special characters. PR1580003
- The IS-IS packet might be corrupted on the provider edge device over the l2circuit tunnel. PR1580047
- On PTX5000 router, the packets might be dropped by the Packet Forwarding Engine after changing the queue of IEEE-802.1ad classifier on FPC-PTX-P1-A or FPC2-PTX-P1A. PR1584042
- There might be higher latency in traffic flow than the configured or default value. PR1588514
- The jsd process might crash in a rare condition in a telemetry scenario. PR1589103
- On PTX3000 and PTX5000 routers, the 40G or 100G interfaces might get become nonresponsive in the `Down` state after the link flaps. PR1589170
- Node name should not be attached to the system hostname under LLDP. PR1593991
- On PTX1000 router, sFlow data (inner VLAN and outer VLAN value, forwarding-class, and DSCP value) does not get exported while checking from the server flow-records at the collector for the ingress sampling. PR1598263
- The CRC errors increases continuously after the interface flaps. PR1600768
- Traffic might silently discarded due to the `RS Fatal` error message on FPC-PTX-P1-A/FPC2-PTX-P1A/FPC-SFF-PTX-P1-A/FPC-SFF-PTX-T. PR1600935
- The l2circuit packets with the destination MAC 01:00:0c:cc:cc:cd might get punted. PR1601360

Layer 2 Ethernet Services

- Copying of files to the RCB over WAN ports is slow. PR1496895

MPLS

- Traffic sent over an LSP might be dropped if two consecutive PLRs along the LSP performs the local repair and bypass protecting the second PLR fails. PR1566101
- Sub-optimal routing issues might occur in case of the LDP route with multiple next-hops. PR1582037
- The LDP replication session might not get synchronized when you enable the dual-transport.
- VPLS connection might get down if you configure the `dual-transport` statement. PR1601854

Multicast

- Multicast traffic in an MVPN setup might be silently discarded on some PTX Series devices that acts as the transit LSR. PR1555274
- FPC might crash in a multicast scenario. PR1569957

Platform and Infrastructure

- FPC might crash in a scaled-firewall configuration. PR1586817
- Upon the receipt of specific sequences of genuine packets destined to the device the kernel crashes and restarts. PR1557881

Routing Protocols

- Traffic might be silently discarded when a BGP route that is part of multipath gets deleted. PR1514966
- Route validation states might flip between the `Valid`, `Invalid`, and `Unknown` state in some corner cases. PR1556656
- The ppmd process memory leaks that might cause traffic loss. PR1561850
- The BGP session carrying the VPNv4 prefix with IPv6 next-hop might be dropped. PR1580578
- The rpd process might crash in certain IS-IS scenario. PR1583484
- The BGP Egress-TE routes loses to the BGP routes using the same protocol-preference. PR1593332

VPNs

- The rpd process might crash during a race condition under the BGP multipath scenario. PR1567918

Resolved Issues: 21.1R1

IN THIS SECTION

Forwarding and Sampling

- The l2ald process might crash due to a next-hop issue in the EVPN-MPLS. PR1548124

General Routing

- On PTX10016 routers, flow control is disabled by default on both aggregated Ethernet interfaces. PR1478715
- In IP-in-IP, end-to-end (CE device to CE device) traceroute is not working as expected. PR1488379
- The following error message might be seen after links flap: `t6e_dfe_tuning_state:et-6/0/0 - Failed to dfe tuning count 10`. PR1512919
- The FPC-E might get stuck. PR1519673
- The chassisd memory leak might cause traffic loss. PR1537194
- Aggregated Ethernet interface framing errors might display increasing values before restoring correct value. PR1539537
- The error message `expr_dfw_action_topo_connect_anh:1434 expr_dfw_action_topo_connect_anh:eda_anh_discard is FALSE for nh-id 568 - return` is observed in PTX1000 routers. PR1540064

- The Packet Forwarding Engine might crash in an MPLS IPv6-tunneling scenario when the next hop changes. PR1540793
- The rpd crash might be seen when BGP service route is resolved over color-only SR-TE policy. PR1550736
- The interface filter with source-port 0 matches everything instead of port 0. PR1551305
- The `show system health-monitor` command is disabled on PTX10008 routers. PR1560268
- On PTX10008 router, BGP next-hop index (indirect and unilist) change after GRES and NSR trigger causes a momentary (unexpected) traffic loss. PR1560323
- The `set chassis display` command is disabled on PTX1008 routers. PR1560453
- An enhancement to enable watchdog petting log on the PTX10000 line cards. PR1561980
- On PTX1000, DCPFE crashes in steady state and generates the following error: `PFE_ERROR_NO_RESOURCE: NH: Failed to alloc for element list`. PR1564147

Infrastructure

- Interface drop counters might display 0 during a race condition when VOQ statistics are also polled simultaneously. PR1537960
- The kernel crashes and generates a core file if churn happens for a flood composite next hop. PR1548545

Interfaces and Chassis

- EOAM IEEE802.3ah link discovery state is Down instead of Active Send Local after deactivating interfaces on routers. PR1532979
- Logs are not being written in /var/log/messages on certain PTX Series platforms. PR1551374

Layer 2 Ethernet Services

- The copying of files to the RCB over WAN ports is slow. PR1496895

MPLS

- Traffic loss might be observed due to rpd crash in an MPLS scenario. PR1528460

Network Management and Monitoring

- A memory leak in the mib2d and snmpd processes might result in SNMP being unresponsive to SNMP queries. PR1543508
- The syslog messages might not be sent with the correct port. PR1545829

Platform and Infrastructure

- The BGP session replication might fail to start after the session crashes on the backup Routing Engine. PR1552603

Routing Policy and Firewall Filters

- Generate route goes to hidden state when the protect core statement is enabled. PR1562867

Routing Protocols

- Traffic might be silently discarded when the clear bgp neighbor all command is executed on a router and also on the corresponding route reflector in succession. PR1514966
- The rpd process generates a core file at gp_rtarget_tsi_update,bgp_rtarget_flash_rt,bgp_rtarget_flash. PR1541768
- BGP-LU session might flap with AIGP scenario. PR1558102

Documentation Updates

There are no corrections or changes in Junos OS Release 21.1R3 documentation for PTX Series routers.

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 21.1

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

NOTE: Back up the file system and the currently active Junos OS configuration before upgrading Junos OS. This allows you to recover to a known, stable environment if the upgrade is unsuccessful. Issue the following command:

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

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

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

To download and install Junos OS Release 21.1R3:

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

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

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

3. Select the 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 composed of Armenia, Belarus, Kazakhstan, Kyrgyzstan, and Russia) can use the following package:

```
user@host> request system software add validate reboot source/junos-install-ptx-
x86-64-21.1R3.9.tgz
```

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

```
user@host> request system software add validate reboot source/junos-install-ptx-
x86-64-21.1R3.9-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 21.1 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 and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

  Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 16: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

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.

Junos OS Release Notes for QFX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for QFX Series switches.

What's New in 21.1R3

There are no new features introduced in this release for QFX Series switches.

What's New in 21.1R2

There are no new features or enhancements to existing features in this release for QFX Series switches.

What's New in 21.1R1

IN THIS SECTION

Software Installation and Upgrade | 286

Learn about new features or enhancements to existing features in this release for QFX Series switches.

Hardware

- **Support for JNP-100G-DAC-1M, JNP-100G-DAC-3M, and JNP-100G-DAC-5M DACs (QFX10002-60C)**—Starting in Junos OS Release 21.1R1, the QFX10002-60C switches support the JNP-100G-DAC-1M, JNP-100G-DAC-3M, and JNP-100G-DAC-5M direct attach copper (DAC) cables.

  [See Hardware Compatibility Tool.]

- **Support for the JNP-QSFP-100G-BXSR and the JNP-QSFP-40G-BXSR bidirectional transceivers**—Starting in Junos OS Release 21.1R1, the QFX5210-64C switches support the JNP-QSFP-100G-BXSR and JNP-QSFP-40G-BXSR bidirectional transceivers.

  [See Hardware Compatibility Tool.]

Authentication and Access Control

- **802.1X authentication on trunk ports (QFX5100 switches)**—Starting in Junos OS Release 21.1R1, you can enable 802.1X authentication on trunk ports on QFX5100 switches. Authentication on the trunk port is supported only in single supplicant and single-secure supplicant modes.

  [See 802.1X Authentication.]

EVPN

- **Tunnel endpoint in the PMSI tunnel attribute field for EVPN Type 3 routes (ACX5448, EX4600, EX4650, EX9200, and QFX10002)**—Starting in Junos OS Release 21.1R1, you can set the tunnel endpoint in the provider multicast service interface (PMSI) tunnel attribute field to use the ingress router's secondary loopback address. When you configure multiple loopback IP addresses on the local provider edge (PE) router and the primary router ID is not part of the MPLS network, the remote PE router cannot set up a PMSI tunnel route back to the ingress router.

  To configure the router to use a secondary IP address that is part of the MPLS network, include the `pmsi-tunnel-endpoint`*`pmsi-tunnel-endpoint`* statement at the `[edit routing-instances` *`routing-instance-name`* `protocols evpn]` hierarchy level for both EVPN and virtual-switch instance types.

  [See EVPN.]

- **Support for remote port mirroring based on VNI match conditions (QFX10002, QFX10008, QFX10016)**—Starting in Junos OS Release 21.1R1, You can use VXLAN network identifier (VNI) values as a match condition when filtering traffic for remote port mirroring. VNI packets that match the configured VNI will be mirrored, with the VNI packet contents, on the designated interface. This addition extends functionality introduced in previous releases.

  [See Filter-based forwarding in EVPN-VXLAN networks and Remote port mirroring to an IP address.]

- **Explicit congestion notification (ECN) over VXLAN tunnels (EX4650 and QFX5120)**—Starting in Junos OS Release 21.1R1, by default, standalone EX4650 and QFX5120 switches support explicit congestion notification (ECN) for packets that are encapsulated across VXLAN tunnels, as follows:

  - During VXLAN encapsulation at the source virtual tunnel endpoint (VTEP), the switch copies the ECN bits of the Type-of-Service (ToS) field from the original packet IP header to the outer VXLAN encapsulation IP header.

  - During VXLAN de-encapsulation at the remote VTEP, the switch copies the ECN bits of the ToS field from the outer VXLAN encapsulation IP header to the original packet IP header.

  You can configure the `vxlan-disable-copy-tos-encap` statement or the `vxlan-disable-copy-tos-decap` statement at the `[edit forwarding-options]` hierarchy on the encapsulation or de-encapsulation ends of the tunnel, respectively, to disable the ECN copy operation.

  > **NOTE**: These switches also copy the differentiated services code point (DSCP) bits in the ToS field of the IP header upon VXLAN encapsulation and de-encapsulation by default, and the same statements disable copying both the DSCP and ECN bits.

  [See vxlan-disable-copy-tos-encap and vxlan-disable-copy-tos-decap.]

Interfaces

- **Dual-speed support on 100Gbps DAC breakout cable (QFX5120-48Y, QFX5200-32C, and QFX5210-64C)**—Starting in Junos OS Release 21.1R1, we support 4x10Gbps speed along with 4x25Gbps speed on the QSFP28 100Gbps DAC breakout cable with the other end SFP28 transceivers. Supported cable lengths are 1, 2, 3, and 5 meters. You can set the 4x10Gbps speed by using the `set chassis fpc` *`fpc`* `pic` *`pic`* `port` *`port number`* `channel-speed 10g` command.

  [See Hardware Compatibility Tool.]

IP Tunneling

- **Support for IPv4 and IPv6 unicast IP-over-IP tunneling (QFX5000)**—Starting in Junos OS Release 21.1R1, we support IP-over-IP tunneling for IPv4 and IPv6 traffic on QFX5000. QFX5000 switches also support recursive route resolution for IP-over-IP tunnels.

  [See Overview of Next-Hop-Based Dynamic Tunneling Using IP-Over-IP Encapsulation.]

- **Support for BGPoverBGPoverBGP recursive route resolution (QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210)**—Starting in Junos OS Release 21.1R1, you can enable BGPoverBGPoverBGP route resolution by creating an expanded route hierarchy, using the `preserve-nexthop-hierarchy` statement at the `[edit routing-options resolution]` hierarchy level.

  [See resolution.]

Junos Telemetry Interface

- **Packet Forwarding Engine and Routing Engine sensor support with JTI (QFX5210)**—Starting in Junos OS Release 21.1R1, you can use Junos telemetry interface (JTI) with remote procedure call (gRPC) services to export Packet Forwarding Engine statistics and Routing Engine statistics from QFX5210 switches to an outside collector. These statistics can also be exported through UDP (native) sensors.

  The supported Packet Forwarding Engine sensors are:

  - Sensor for CPU (microkernel) memory (resource path **/junos/system/linecard/cpu/memory/**)
  - Sensor for firewall filter statistics (resource path **/junos/system/linecard/firewall/**)
  - Sensor for physical interface traffic (resource path **/junos/system/linecard/interface/**)
  - Sensor for logical interface traffic (resource path **/junos/system/linecard/interface/logical/usage/**)
  - Sensor for software-polled queue-monitoring statistics (resource path **/junos/system/linecard/qmon-sw/**)

  The supported Routing Engine sensors are:

  - Sensor for LACP state export (resource path **/lacp/**)
  - Sensor for chassis environmentals export (resource path **/junos/system/components/component/**)
  - Sensor for chassis components export (resource path **/components/**)
  - Sensor for LLDP statistics export (resource path **/lldp/interfaces/interface[name='*name*']/**)

- Sensor for BGP peer information export (resource path **/network-instances/networkinstance/protocols/protocol/bgp/**)
- Sensor for RPD task memory utilization export (resource path **/junos/task-memoryinformation/**)
- Sensor network discovery ARP table state (resource path **/arp-information/**)
- Sensor for network discovery NDP table state (resource path **/nd6-information/**)

[See Understanding OpenConfig and gRPC on Junos Telemetry Interface.]

Layer 2 Features

- **Configurable EtherType values (QFX10002-36Q, QFX10002-72Q, and QFX10008)**—Starting in Junos OS Release 21.1R1, you can customize the EtherType field values stored in the ternary content addressable memory (TCAM) tables. Ethernet frame headers contain an EtherType field to identify the protocol in the frame's payload so the receiving device knows how to process the traffic. The device keeps a default list of the EtherType values it can process in TCAM for fast access. With this feature, you can define custom EtherType values in place of some of the default values in the TCAM table either for a specified FPC slot or for all currently active FPCs on the switch. Some EtherType values are reserved; you can't change or reconfigure those values.

  [See ether-type.]

MPLS

- **Nonstop active routing (NSR) support for controller-initiated RSVP label-switched paths (LSPs) (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 21.1R1, we support NSR for controller-initiated RSVP-based point-to-point (P2P) and point-to-multipoint (P2MP) LSPs. The primary Routing Engine synchronizes all RSVP LSPs initiated by Path Computation Elements (PCEs), including multicast flow specifications for any PCE-initiated P2MP LSPs, with the backup Routing Engine. This ensures zero traffic loss for the traffic carried over PCE-initiated RSVP LSPs during Routing Engine switchovers. This feature is enabled when NSR is configured.

  [See PCEP Configuration.]

Multicast

- **Controller-based BGP multicast signaling (MX Series, PTX Series, and QFX Series)**—Starting in Junos OS Release 21.1R1, we've introduced controller-based BGP multicast signaling that can be used instead of hop-by-hop signaling to program multicast forwarding states on routers. An external controller that is aware of the topology and network events within that topology calculates the optimum multicast trees between the source and receivers. The external controller then uses BGP signaling to send a new type of BGP network layer reachability information (NLRI) with modified attributes to convey the multicast state information to all the routers on the multicast trees.

You can use this feature instead of multicast routing protocols, such as Protocol Independent Multicast (PIM) or multipoint LDP (MLDP). You can enable this feature using `bgpmcast` configuration option at the `[edit protocols]` hierarchy.

- **Support for next-generation multicast VPN (QFX10002, QFX10008, and QFX10016)**—Starting in Junos OS Release 21.1R1, we support Multiprotocol BGP (MBGP) next-generation multicast VPNs with the following types of provider tunnels:

  - Ingress replication
  - RSVP-Traffic Engineering (RSVP-TE) point-to-multipoint (P2MP)
  - Multipoint LDP P2MP

  A P2MP is a Multiprotocol Label Switching (MPLS) label-switched path (LSP) with a single source and multiple destinations. By taking advantage of MPLS packet replication capability of the network, P2MP LSPs avoid unnecessary packet replication at the ingress router. Packet replication takes place only when packets are forwarded to two or more different destinations requiring different network paths.

  [See Multiprotocol BGP MVPNs Overview and provider-tunnel.]

Network Management and Monitoring

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

  [See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

- **sFlow support for IP-IP traffic (PTX1000, PTX10008, and QFX10002)**—Starting in Junos OS Release 21.1R1, you can use sFlow technology to sample IP over IP (IP-IP) traffic on a physical port. sFlow sampling is supported for IP-IP tunnels that have an IPv4 outer header that carry IPv4 or IPv6 traffic. You can use sFlow monitoring technology to randomly sample network packets from IP-IP tunnels

and to send the samples to a destination collector for monitoring. Devices that act as an IP-IP tunnel entry point, transit device, or tunnel endpoint support sFlow sampling.

[See Overview of sFlow Technology and Configuring IP Tunnel Interfaces.]

- **Remote port mirroring to IPv6 address (GRE encapsulation)(EX4650, EX4650-48Y-VC, QFX5120, QFX5120-32C, QFX511120-48T, QFX5120-48T-VC, QFX5120-48Y, and QFX5120-48YM)**—Starting in Junos OS Release 21.1R1, you can use remote port mirroring to copy packets entering a port or VLAN and sends the copies to the IPv6 address of a device running an analyzer application on a remote network (sometimes referred to as “extended port mirroring”). When you use remote port mirroring the mirrored packets are GRE-encapsulated.

  Add the address you would like to have the copied packets sent to in the CLI hierarchy. For example, `set forwarding-options analyzer ff output ipv6-address 2000::1`.

  [See Understanding Port Mirroring and Analyzers.]

Routing Policy and Firewall Filters

- **Support for microsegmentation on VLANs and VXLANs (QFX5110 and QFX5120)**—Starting in Junos OS Release 21.1R1, you can configure egress filters with Layer 2 and Layer 3 match conditions in both VLAN and VXLAN deployments. Junos OS already supports filtering in Layer 2 match conditions in the ingress direction.

  To use egress filters for microsegmentation in a VXLAN, enable the `epacl-firewall-optimization` statement at the `[edit chassis]` level of the hierarchy and create the firewall rules with the match conditions that you want to filter on. For egress filtering on VLANs, you don't need to enable `epacl-firewall-optimization`. Both the QFX5110 and QFX5120 support egress filtering, for VLANs and VXLANs, with the following match conditions:

  - `ip-source-address`
  - `ip-destination-address`
  - `destination-port`
  - `destination-mac-address`
  - `user-vlan-id`
  - `ip-protocol`
  - `source-mac-address`

  Valid actions for these rules are `accept`, `count`, and `discard`.

  [See Overview of Firewall Filters (QFX Series) and Understanding Firewall Filter Match Conditions.]

Software Installation and Upgrade

- **Zero-touch provisioning (ZTP) with IPv6 support (QFX5120-32C)**—Starting in Junos OS Release 21.1R1, you can use a DHCPv6 client and ZTP to provision a QFX5120-32C switch.. During the bootstrap process, the device first uses the DHCPv4 client to request for information regarding the image and configuration file from the DHCP server. The device checks the DHCPv4 bindings sequentially. If there is a failure with one of the DHCPv4 bindings, the device continues to check for bindings until provisioning is successful. However, if there are no DHCPv4 bindings, the device checks for DHCPv6 bindings and follows the same process as for DHCPv4 until the device is provisioned successfully. Both DHCPv4 and DHCPv6 clients are included as part of the default configuration on the device.

  The DHCP server uses DHCPv6 options 59 and 17 and applicable suboptions to exchange ZTP-related information between itself and the DHCP client.

  [See Zero Touch Provisioning.]

- **Support for signed third-party application installation (MX10003, MX10008, QFX5210, QFX10002, and QFX10008 routers with VM host architecture)**—Starting in Junos OS Release 21.1R1, you can install signed third-party application installation and carry over the application between upgrades.

  The backup of third-party package occurs during upgrade. Hence, the package is restored even if the installed package is deleted or uninstalled before a reboot. However, as the third party package restoration depends on the contents saved on the disk during upgrade and the configuration to allow the package to be installed, restoration is not possible when

  - Configuration is removed after upgrade
  - Content is removed due to deletion by configuring `request vmost zeroize` command

  On platforms where jinstall-host.tgz images are installed, the minimum space required for the backup is 250MB. After backup, if the free space available is less than 200MB, the backup would be deleted to make space for upgrade. On platforms where junos-vmhost images are installed, the minimum space required for backup of third party unbundled packages is 1200MB. After the backup, if the free space is less than 512MB, the backup would be deleted to free up space for upgrade.

  [See Installing, Upgrading, Backing Up, and Recovery of VM Host.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for QFX Series Switches.

What's Changed in Release 21.1R3

IN THIS SECTION

EVPN

- **Community information no longer included in VRF routing table**— The QFX series switches will no longer include the inherited advertised route target communities, EVPN extended communities, or vxlan encapsulation communities for EVPN Type 2 and EVPN Type 5 routes when an IP host is added in the VRF routing table.

Interfaces and Chassis

- When configuring multiple flexible tunnel interface (FTI) tunnels, the source and destination address pair needs to be unique only among the FTI tunnels of the same tunnel encapsulation type. Prior to this PR, the source and destination address pair had to be unique among all the FTI tunnels regardless of the tunnel encapsulation type.

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

  [See request system scripts refresh-from and cert-file (Scripts).]

Layer 2 Features

- **Configurable EtherType values (QFX10016)**—You can now customize the EtherType field values stored in the ternary content addressable memory (TCAM) tables. With this feature, you can define custom EtherType values in place of some of the default values in the TCAM table either for a specified FPC slot or for all currently active FPCs on the switch. Some EtherType values are reserved; you can't change or reconfigure those values.

What's Changed in Release 21.1R2

IN THIS SECTION

General Routing

- **SSH session connection limit and rate limit per connection (PTX Series and QFX Series)**—We have introduced SSH `connection-limit` and `rate-limit` options at the `edit system services ssh` hierarchy levels to

enable SSH connection limit and rate limit per connection. The default connection limit value is 75 connections and there is no default value associated with rate limit.

- **Support only for manual channelization on QSFP-100G-SR4-T2 optics (QFX5120-48T and QFX5120-32C)**—We recommend that you use the active optical cable (AOC) for auto-channelization. The QSFP-100G-SR4-T2 cables do not support auto-channelization. To use the QSFP-100G-SR4-T2 optics with an external breakout cable, you must configure the channelization manually by running the `channel-speed` statement at the **edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)** hierarchy level.

  [See channel-speed.]

- **Juniper Agile Licensing (QFX5120-48Y, QFX5110-32Q, and QFX5110-48S)—**Starting from this release onwards, the QFX switch supports following features:

  - **Standard:**BFD, Filters (Layer 2 and Layer 3), Layer 2 (xSTP, 802.1Q, LAG), Layer 3 (static), QoS (Layer 2 and Layer 3), and SNMP
  - **Advanced 1:** Standard features, BGP, IS-IS, FBF, VRRP, MC-LAG, Layer 3 (static), GRE tunnel, OSPF, RIP, sFlow, and Virtual Chassis
  - **Advanced 2:** Advanced 1 features, CFM, Q-in-Q, VXLAN, PCEP, ESI-LAG, Timing, Ethernet OAM, EVPN-VXLAN, IGMP version 1, IGMP version 2, and IGMP version 3, PIM, and Multicast Listener Discovery (MLD) version 1 or version 2
  - **Premium:** Advanced 2 features, Layer 3 VPN, LDP, RSVP, Layer 2 circuit, EVPN-MPLS, Segment routing, MPLS, and MACsec

  [See Flex Software License for QFX Series Switches and Juniper Agile Licensing Guide.]

Interfaces and Chassis

- **Blocking duplicate IP detection in the same routing instance (All Junos platforms)**—Junos will no longer accept duplicate IPs between different logical interfaces in the same routing instance. Refer to the table mentioned in the topic inet (interfaces). When you try to configure same IP on two logical interfaces inside same routing instance, the commit will be blocked with the error displayed as shown below: **edit** user@host# set interfaces ge-0/0/1 unit 0 family inet address 2.2.2.2/24 **edit** user@host# commit commit complete **edit** user@host# set interfaces ge-0/0/2 unit 0 family inet address 2.2.2.2/24 **edit** user@host# commit **edit interfaces ge-0/0/2 unit 0 family inet** 'address 2.2.2.2/24' identical local address found on rt_inst **default**, intfs **ge-0/0/2.0 and ge-0/0/1.0**, family **inet**. error: configuration check-out failed

  [See inet(interfaces).]

Layer 2 Ethernet Services

- **Link selection support for DHCP (QFX Series)**—We have introduced the `link-selection` statement at the **edit forwarding-options dhcp-relay relay-option-82** hierarchy level, which allows DHCP relay to add suboption 5 to option 82. Suboption 5 allows DHCP proxy clients and relay agents to request an IP address for a specific subnet from a specific IP address range and scope. Prior to this release, the DHCP relay dropped packets during the renewal DHCP process and the DHCP server used the leaf's address as a destination to acknowledge the DHCP renewal message.

  [See relay-option-82.]

Network Management and Monitoring

- `Change in OID ifHighSpeed`—Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

  [See SNMP MIBs and Traps Supported by Junos OS.]

- **Enhancement to the snmp mib walk command (PTX Series, QFX Series, EX Series, MX Series, SRX Series)** —The ipv6IfOperStatus field displays the current operational state of the interface. The noIfIdentifier(3) state indicates that no valid Interface Identifier is assigned to the interface. This state usually indicates that the link-local interface address failed Duplicate Address Detection. When you specify the 'Duplicate Address Detected' error flag on the interface, the new value (noIfIdentifier(3)) is displayed. Previously, the snmp mib walk command did not display the new value (noIfIdIdentifier(3)).

- **Changes in contextEngineID for SNMPv3 INFORMS (PTX Series, QFX Series, ACX Series, EX Series, MX Series, and SRX Series**—Now the contextEngineID of SNMPv3 INFORMS is set to the local engine-id of Junos devices. In earlier releases, the contextEngineID of SNMPv3 INFORMS was set to remote engine-id.

  See SNMP MIBs and Traps Supported by Junos OS.

What's Changed in Release 21.1R1

IN THIS SECTION

General Routing

- **SSH session connection limit and rate limit per connection (PTX Series and QFX Series)**—We have introduced SSH `connection-limit` and `rate-limit` options at the `edit system services ssh` hierarchy levels to enable SSH connection limit and rate limit per connection. The default connection limit value is 75 connections and there is no default value associated with rate limit.

- **Change in license bandwidth command on vMX virtual routers**—Starting in Junos OS 21.1R1, to use the available license bandwidth, explicitly set the license bandwidth use the `set chassis license bandwidth <in Mbps>` command.

  [See Configuring Licenses on vMX Virtual Routers..]

- **Support only for manual channelization on QSFP-100G-SR4-T2 optics (QFX5120-48T and QFX5120-32C)**—We recommend that you use the active optical cable (AOC) for auto-channelization. The QSFP-100G-SR4-T2 cables do not support auto-channelization. To use the QSFP-100G-SR4-T2 optics with an external breakout cable, you must configure the channelization manually by including the `channel-speed` statement at the **edit chassis fpc slot-number pic pic-number (port port-number | port-range port-range-low port-range-high)** hierarchy level.

  [See channel-speed.]

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX event scripts. If you include the parameter, the function does not generate and log UI_LOGIN_EVENT and

UI_LOGOUT_EVENT messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root UI_LOGIN_EVENT and UI_LOGOUT_EVENT messages are included in system log files.

[See invoke() Function (SLAX and XSLT).]

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Layer 2 Ethernet Services

- **Modification to sync-reset command (All JUNOS and EVO platforms)**—Starting from this release, the sync-reset command is disabled by default on all Junos and EVO platforms. Sync-reset command enables the device to send the sync bit in the LACP packets on minimum-link failure. Previously the sync-reset command was enabled by default on QFX and EX series, while it was by default disabled on MX, PTX and ACX series.

  [See sync-reset.]

Network Management and Monitoring

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the `[edit system services netconf hello-message yang-module-capabilities]` hierarchy level. In addition, you can specify the YANG schemas that the NETCONF server should include in its list of supported schemas by

configuring the appropriate statements at the `[edit system services netconf netconf-monitoring netconf-state-schemas]` hierarchy level.

[See hello-message and netconf-monitoring.]

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the `verbose` statement at the `[edit system export-format json]` hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from `verbose` to `ietf` starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `[edit system export-format json]` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

  [See export-format.]

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for QFX Series switches.

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

General Routing

- Junos OS can hang trying to acquire the SMP IPI lock while rebooting when it is running as a VM on Linux and QEMU hypervisor. Device can be recovered using power-cycle of the device. PR1385970

- This issue occurs due to PECHIP limitation when underlay is tagged. After de-encapsulation when inner packet is recirculated it still retains the VLAN tag property from the outer header since outer header was tagged. Thus 4 bytes of inner tag got overwritten in inner packet and packet got corrupted which will result in EGP chksum trap seen in PECHIP. Fixing PECHIP limitation in software has high risk. It will be accommodated in a future release. Workaround is provided to enable `encapsulate-inner-vlan` configuration statement. PR1435864

- This warning will pop up at reboot or power off time. This notification is part of unmount routine which is harmless with no functional impact. This might get fixed in coming RPCL/WRL release. PR1527581

- In case of fan failure, `show chassis environment` and `show chassis fan` will show `failed` and `check` status respectively. This is expected and no discrepancy in terms of real status. PR1527628

- On QFX10002 switches in a dynamic IP-IP tunnel transit scenario, when sFlow egress sampling is enabled on an aggregated Ethernet interface in an ECMP case, the sFlow export data does not include the nextHop field. PR1533307

- When an image with the third party SDK upgrade (6.5.x) is installed, the CPU utilization might go up by around 5 percent. PR1534234

- ECMP over GRE does not work for BGP routes. Traffic is polarized to just one egress interface but not distributed to multiple egress interfaces. PR1537924

- In QFX5100 and EX4300 non-TVP platforms, the sample rate is limited by the IPC between the Packet Forwarding Engine and the sFlow process, so the supported limit is around 700 samples per second in these platforms. This is applicable to any sampled packets in these platforms and not specific to IPnIP. PR1539815

- ISSU is not supported from releases below 20.4 to releases 20.4 and above. There is a major SDK upgrade from 6.3.2 to 6.5.16, due to which the warm boot feature needed for ISSU is not supported by our vendor. PR1554915

- On QFX5200 and QFX5100 switches with the IPIP tunnel feature, `show dynamic-tunnels database statistics` command output shows extra packet counts (i.e. sampled packets when sFlow is enabled). PR1555922

Infrastructure

- Software versions 21.1 and lower are running FreeBSD version 11 whereas from version 21.2 onward, the FreeBSD version is 12. Software upgrade to 21.2 (or later) from 21.1 (or earlier) will mandatorily need cli configuration statement `no-validate` to be used during software image upgrade process. (For EX4400 platforms, this is applicable from version 21.3 onward. Hence, for EX4400

platforms, software upgrade to 21.3 (or later) from 21.2 (or earlier) will mandatorily need cli configuration statement `no-validate` to be used during software image upgrade process.) PR1586481

Routing Protocols

- On QFX-5210 platforms, when two Flex Hash rules are configured, on deactivating first one, second one is not programmed in hardware. Commit works though. 1. Two flex hash profiles with same traffic type and different hash-parameters cannot be configured. For example: * Profile 1: Below profile applies flex hash rule for traffic with mpls 2 -label and pick the offsets configured set forwarding-options enhanced-hash-key flex-hashing FH-MPLS-2-V4-TCP-UDP ethtype mpls num-labels 2 set forwarding-options enhanced-hash-key flex-hashing FH-MPLS-2-V4-TCP-UDP ethtype mpls conditional-match CM-MPLS-2-V4-TCP set forwarding-options enhanced-hash-key flex-hashing FH-MPLS-2-V4-TCP-UDP ethtype mpls hash-offset offset1 base-offset1 start-of-L3-OuterHeader set forwarding-options enhanced-hash-key flex-hashing FH-MPLS-2-V4-TCP-UDP ethtype mpls hash-offset offset1 offset1-value 28 set forwarding-options enhanced-hash-key flex-hashing FH-MPLS-2-V4-TCP-UDP ethtype mpls hash-offset offset1 offset1-mask ffff * Profile 2: Below profile applies flex hash rule for traffic with mpls 2-label which is same as profile 1. Since already there is a profile1 configured to match mpls label-2 traffic this profile will not be installed and doesn't have any impact . set forwarding-options enhanced-hash-key flex-hashing FH1-MPLS-2-V4 ethtype mpls num-labels 2 set forwarding-options enhanced-hash-key flex-hashing FH1-MPLS-2-V4 ethtype mpls hash-offset offset1 base-offset1 start-of-L3-OuterHeader set forwarding-options enhanced-hash-key flex-hashing FH1-MPLS-2-V4 ethtype mpls hash-offset offset1 offset1-value 0 set forwarding-options enhanced-hash-key flex-hashing FH1-MPLS-2-V4 ethtype mpls hash-offset offset1 offset1-mask ffff 2. Two flex hash profiles with same traffic types can be configured only when below conditions are met : 1. All the flex-hash parameters like base-offset, offset-value, offset-mask are same for both profiles. 2. Both profiles should have different conditional profiles attached. 3. The conditional parameters like base-offset and offset-value have to be same for both profiles. 4. Combination of **match-data & match-mask** has to be different for both profiles . PR1521306

Open Issues

IN THIS SECTION

Learn about open issues in this release for QFX Series switches.

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 QFX10000, source MAC and TTL values are not updated for routed multicast packets in EVPN-VXLAN. PR1346894
- Backup Routing Engine might crash after GRES occurs continuously for more than 10 times. PR1348806
- When VLAN is added as an action for changing the VLAN in both ingress and egress filters, the filter is not installed. PR1362609
- Due to transient hardware condition, single-bit error (SBE) events are corrected and have no operational impact. Reporting of those events had been disabled to prevent alarms and possibly unnecessary hardware replacements. This change applies to all Platforms using Hybrid Memory Controller (HMC). PR1384435
- Storm-control does not rate-limit ARP packets on QFX10000 although shutdown action works. PR1461958
- VXLAN VNI (multicast learning) scaling on QFX5110 traffic issue is seen from VXLAN tunnel to Layer 2 interface. PR1462548
- On QFX5000 series platforms with "instance-import", deleting route which has "next-table" used might result in unexpected route next-hop. PR1477603
- When running the command, `show pfe filter hw filter-name filter name`, the command fails to retrieve the Packet Forwarding Engine programming details of the filter. PR1495712

- On Junos QFX platform, the Jflow service might not report the accurate throughput rate. This issue is seen when there is high sampled traffic rate with low flow cache hit ratio. PR1502645
- After repeated deletion and addition of logical switch on NSX-V setup, with OVSDB configured, ping between the VM to the baremetal server fails intermittently (only on a few iterations out of the total number of iterations). PR1506097
- Changing the scaled firewall profiles runtime (on the fly) does not release the TCAM resources as expected. PR1512242
- On the QFX10000 line of switches, when an explicit Layer 2 classifier is applied on a Layer 3 interface, the default Layer 3 classifiers are not removed. By design, the Layer 3 classifier takes precedence over the Layer 2 classifier. PR1520570
- The MSDP sessions might reset after a GRES reset even when nonstop routing state is synchronized and ready for switchover. PR1526679
- FIPS mode is not supported. PR1530951
- On QFX5100 devices that do not run the QFX-5e codes (non-TVP architecture), when an image with the vendor SDK upgrade (6.5.x) is installed, the CPU utilization may go up by around 5%. PR1534234
- On QFX10002 devices acting as PHP, egress sFlow samples do not report MPLS explicit-null label in the raw packet header. The MPLS payload can be of IPv4 or IPv6 protocol. PR1537946
- **Socket to sflowd closed** error is seen when the ukern socket to sflowd daemon (server) is closed. The error is rectified by itself when the connection is re-established in the subsequent attempts. When these errors are consistent, it indicates the communication issue between sFlow running on the FPC with the sflowed. PR1538863
- EVPN-VXLAN: vmcore seen on primary and backup Routing Engines of QFX10008 with Layer 2 or Layer 3 multicast configuration. PR1539259
- The BCMX calls are deprecated and needs to be replaced with BCM calls. PR1541159
- Moving WRL7 SDK to RCPL31 for QFX10000 platforms. RCPL31 provides tweaking of build infra and fixes related to performance enhancements and vulnerabilities in WRL7 Linux. PR1547565
- 100G AOC from Innolight does not come up after multiple reboots. It recovers after the interface is disabled and then enabled. PR1548525
- 5M DAC connected between QFX10002-60C and MX2010 doesn't link up. But with 1M and 3M DAC, this interoperation works as expected. Also it is to be noted on QFX10002-60C and ACX Series Devices or traffic generator, the same 5M DAC works seamlessly. There seems to be a certain SI or link-level configuration on both QFX10002-60C and MX2010 which needs to be debugged with the help from HW and SI teams and resolved. PR1555955

- In Junos OS Release 20.2, some features show up as a licensed features. Customer might see alarms, commit warnings, and the following `show system license`. However, there would be no functional impact. user@router> show system license License usage: Licenses Licenses Licenses Expiry Feature name used installed needed esi-lag 1 0 1 invalid.PR1558017
- To avoid the additional interface flap, interface hold time needs to be configured. PR1562857
- Starting 21.1R1, Junos will be shipping with python3 (python2 is no longer supported). In ZTP process, if a python script is being downloaded, please ensure the python script follows python3 syntax (there are certain changes between python2 and python3 syntax). Also, so far (ie until 20.4R1), the python script had #!/usr/bin/python as the first line (ie the path of the python interpreter). The same needs to be changed to #!/usr/bin/python3 from 21.1R1. PR1565069
- The pic_create_ifname: 0/0/0 pic type F050 not supported log messages generated under chassisd and other messages in logs. PR1566440
- In mixed QFX5100, EX4300 VCF setup, duplicate traffic might be observed for some Layer 3 multicast traffic streams . PR1568152
- The Broadcast, Unknown Unicast, and Multicast (BUM) traffic replication over VTEP is sending out more packets than expected and there seems to be a loop also in the topology. PR1570689
-
- On QFX5100, while checking DHCP smart relay over IRB interfaces, the renew-ack's may not be seen in the dhcp client. PR1581025
- On QFX series, switches with the vendor chip as Packet Forwarding Engine, if IS-IS is enabled on an integrated routing and bridging (IRB) interface and the maximum transmission unit (MTU) size of the IRB interface is configured with a value great than 1496 bytes, the IS-IS hello (IIH) PDUs with jumbo frame size (i.e., great than 1496 bytes) might be dropped and not sent to the IS-IS neighbors. The following is the product list of QFX series switches with vendor chip as Packet Forwarding Engine. QFX3500/QFX3600/QFX5100/QFX5110/QFX5120/QFX5130/QFX5200/QFX5210/ QFX5220PR1595823
- QFX10002: dcpfe core is observed after booting the device with EVPN-VXLAN configurations. PR1597479
- Read write lock is not acquired during the sysctl invocation. The assert triggered in the interface state function call leads to go Routing Engine 1 to debug (db>) prompt. PR1598814
- After performing an upgrade, the peer device is rebooted or the peer interface is disabled or enabled, then the SFP-T port might remain in up state but could not forward traffic. PR1600291
- Convergence time degradation is seen in IS-ISv6, OSPFv2, and OSPFv3. PR1602334

- On QFX10000 platforms supporting DHCP (Dynamic Host Configuration Protocol), when relay is configured in non-default routing instance and DHCP renew process is triggered due to lease time expiry, renew packets might get flooded to all members in VLAN instead of unicast forwarding. PR1603444

- On QFX5120, traffic loss might be seen when primary link disabled with aggregated Ethernet Link Protection configuration. PR1604350

- On QFX10008 and QFX10016 platforms, the system reboot takes approximately 9 minutes for FPCs to come online after `system reboot` command is issued. PR1605002

- Dfwd cored when accessing ephemeral db files which is deleted through script. PR1609201

- In QFX10002-60C under MAC statistics `output-mac-control-frames` and `output-mac-pause-frames` does not increment. PR1610745

- As per design change, QFX10002, QFX10008, and QFX10016 devices have max token allocation limit of 104,000 combined for l2/l3/BUM routes. These tokens are used as egress nexthop IDs for route lookup in HW. Out of 104,000 tokens, 8000 tokens are reserved for maximum VXLAN tunnels and 16,000 tokens are reserved for IRB interfaces, and locally learned ARP/NDP nexthops also require tokens to install in HW which will be allocated from free pool of tokens left after the reservations. So, total tokens available in free pool after reserving [8000 (vxlan tunnels) + 16,000(IRBs)] is 80,000 tokens. Thus, maximum of 80,000 local ARP entries can be supported on QFX10,000 (Elit/Ultimate) devices. PR1616224

- On all QFX5000 platforms, when the IRB (Integrated Routing and Bridging) is configured and removed from the VLAN, Layer3 bits enabled in hardware are not reset, hence the Layer 3 lookup on the VLAN is not deactivated while IP address/IRB is not configured. So, any Layer 2 traffic coming with dmac as my mac might be dropped. PR1618425

- Log messages **fpc0 SRIRAM Tx VxLAN Ucast: ifd_out = vtep dst_gport is (c00000X) so do not process pkt further** can show up on QFX5000 switches. These are harmless messages. PR1624925

- On QFX5100 platforms with the Virtual Chassis scenario, if the Virtual Chassis Ports (VCPs) are connected through QSFP+40GE-AOC cable, post upgrading to 17.3 or later releases, VCPs might not come up or flap impacting VC functionality and services. PR1633998

- On QFX platforms in VxLAN scenario, if STP is enabled on all the interfaces of the switch, ethernet table might not get populated to locally connected devices resulting in traffic blackhole. PR1636950

- When L2PT (Layer2 Protocol Tunneling) is enabled on a transit switch using SP style configuration, protocol convergence between end nodes might fail. PR1637249

EVPN

- End-hosts might not communicate via Ethernet VPN with Virtual Extensible LAN encapsulation (EVPN-VxLAN) domain after Ethernet Segment Identifier (ESI) failover. This issue affects QFX5000 platforms only. Please refer to restoration steps when this issue is encountered. PR1584595
- Modifying the I-ESI value is traffic effecting event. If this must be done then follow the below steps in order to avoid this PR 1) deactivate interconnect stanza for the routing-instance in question 2) Modify the I-ESI value 3) activate the interconnect stanza. PR1600600
- In all Junos and Junos OS Evolved platforms, EVPN-VXLAN scenario, with `proxy-macip-advertisement` statement configured, few ARP/ND/MAC entries might get missing. PR1609322
- This problem happens only with the translation VNI when MAC is moved from DC1 to DC2. VM moves across DC where there is no translate VNI configuration in the interconnect works as designed. PR1610432
- Multiple memory leaks might be seen, which might lead to the process rpd crash. Issue 1- On QFX platforms configured with EVPN-VxLAN memory leaks might be seen due to BGP communities. Issue 2- On all Junos platforms memory leaks might be seen due to MAC mobility. PR1626416

Infrastructure

- The following messages are seen during FTP: `ftpd[14105]: bl_init: connect failed for /var/run/blacklistd.sock (No such file or directory)` messages are seen during FTP. PR1315605

Layer 2 Features

- On QFX5100 platforms, if a change related to TPID is made in the Device Control Daemon, traffic might be dropped in Packet Forwarding Engine due to failure on layer 2 learning or interfaces flapping. PR1477156

Layer 2 Ethernet Services

- It was observed occassionally that issuing a **request system zeroize** did not trigger ZTP. A simple workaround is to re-initiate ZTP. PR1529246

- On all Junos platforms configured as DHCP (Dynamic Host Configuration Protocol) server or relay-agent, the file system storage under /var directory might get filled up with DHCP ERA (Event Rate Analyzer) logs which is enabled by default and could result in other processes not having storage space to log details of router functionalities. PR1617695
- With Non-Flex images ZTP will not work on QFX5200. PR1629441

MPLS

- The rsvp interface update threshold configuration syntax has changed between Junos OS Release 18.2X75-D435 and Junos OS Release 20.3X75-D10 to include curly braces around the threshold value. Upgrading and downgrading between these releases is not entirely automatic. PR1554744

Platform and Infrastructure

- Arrival rates were not seen at system level when global-disable fpc is configured on QFX. PR1438367
- When the DHCP relay mode is configured as no-snoop, we are observing the offer gets dropped due to incorrect ASIC programming. PR1530160
- When the DHCP relay mode is configured as no-snoop, we are observing the offer gets dropped due to incorrect ASIC programming. This issue only affects while running DHCP relay on EVPN/VXLAN environment. PR1530160

Routing Policy and Firewall Filters

- On all Junos OS platforms with **set policy-options rtf-prefix-list** configured, when upgraded to a specific version, the device might fail to validate its configuration, which eventually causes rpd to crash unexpectedly. The reason should be a software fault. PR1538172

Routing Protocols

- On QFX5000 platforms, when the host forwarding table is full and the host entries are installed in the LPM forwarding table, or when lpm-profile with unicast-in-lpm option is used, the Layer 3 IP

route might not be installed in the LPM forwarding table if there are SER errors, hence there might be traffic impact. PR1429504

- Currently IPIP, IPv6 and gre decapsulations are supported. It is not recommended to configure gre and IPIP/IPv6 are configured in a single filter, then the last decapsulate filter term is used to program the entire filter terms. PR1580468
- On QFX10002 platforms, the multi-hop BFD session might flap if collecting RSI or some other outputs (such as show interface or configuration). It is caused by the missing BFD packets because the PPMAN thread is not scheduled within the BFD timers which are 300 milliseconds with a multiplier of 3. PR1589765
- On all platforms, traffic drops might be seen when incorrect VPN labels are allotted. When there is a change in the nexthop by BGP policy, the traffic is still forwarded to the old label. This leads to traffic drops for prefixes sending traffic to the old nexthop. PR1617691
- In the single-hop BFD of BGP scenario, when multiple addresses of the same subnet are configured on the interface of the BFD session, the BFD session might be down. PR1635700

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for QFX 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: 21.1R3

IN THIS SECTION

General Routing

-

- The Virtual Chassis Port (VCP) might not come up QFX5100 platform PR1555741
- On the QFX5110 line of switches, the untagged traffic routed over native-vlan might be dropped.PR1560038
- Multiple entries to vlan-id-list might not work in EVPN-VXLAN scenario. PR1564403
- The MAC address will point to incorrect interface after traffic is stopped and not aging out. PR1565624
- BFD sessions over VTEP might fail. PR1571417
- On the QFX10000 line of switches, the dcpfe or fpc process might crash if the ARP MAC occurs. PR1572876
- On QFX10K2-60C platforms, disk missing alarms are not seen. PR1573139
- The OSPF session over IRB might not come up in the EVPN-VXLAN scenario PR1577183
- On the QFX10000 line of switches, the port might not go in to the Down state immediately during some abnormal type of line card reboots.PR1577315
- Kernel crash might be observed on the backup Routing Engine after GRES. PR1577799

- On the QFX5100 line of switches, few 40G ports might not be channelized successfully. PR1582105
- The pciephy and firmware download is not working after migrating to 6.5.19. PR1582244
- The srxpfe process might crash. PR1582989
- Firewall filter is not getting programmed after deleting a large filter and adding a new one in a single commit on QFX5000 platforms.PR1583440
- On QFX5000 series, the `show route detail` might not show Next-hop type IPoIP Chained comp nh in the output (Display only - no operation impact). PR1584322
- The QFX5000 and QFX10000 devices might get hanged for sometime after reboot. PR1584902
- On QFX10002-60C platform, high FPC CPU utilization might be seen. PR1585728
- The firewall filter matching condition **then dscp** is not supported in the VC scenario. PR1586600
- The na-grpc process crash might be seen and existing telemetry connections will be disconnected.PR1587956
- The dcpfe might crash when loading EVPN with VXLAN configuration on the setup. PR1588637
- On the QFX5210-64C line of switches, PSU firmware upgrades through JUNOS. PR1589572
- On the QFX5120 line of switches, the MPLS traffic might not be forwarded after the aggregate interface flaps.PR1589840
- VC primaryship changed and connection dropped after renumbering of backup member ID. PR1590358
- On QFX5120-48T after removing 1g speed on interfaces it won't come back as 10g PR1591038
- PTP Queue might become stuck under specific traffic pattern and GM switchover. PR1591571
- The IPv4 fragmented packets might be broken if PTP transparent clock is configured. PR1592463
- BFD session might flap during Routing Engine switchover. PR1593244
- The dcpfe process might crash in EVPN-VxLAN scenario. PR1593950
- Packet drop might occur in ECMP next-hop flap scenario. PR1594030
- ARP entry might be found missing intermittently post FPC reboot. PR1594255
- The label field for the EVPN Type 1 route is set to 1. PR1594981
- The reinstallation of the Type-5 tunnels might fail in the EVPN-VXLAN scenario. PR1595197

- The DCI InterVNI and IntraVNI traffic might black-holed in gateway node due to the tagged underlay interfaces. PR1596462
- The mcscnoopd process generates the core files at `snp_token_db_gencfg_handler,krt_decode_gencfg,krt_ifstate_resync_read,krt_async_recv_ifstate_resync_phase`. PR1596483
- The l2alm fails to send IPC message to the l2ald which might cause the FPC to crash. PR1596615
- The `fpc0 bcm pkt reinsert failed` log written in the log messages in an aggressive way. PR1596643
- TTraffic might be dropped after backup FPC is rebooted in a Virtual Chassis scenario. PR1596773
- The interface might not be brought up when Q-in-Q is configured.PR1597261
- Deletion of MACsec configuration on a logical interface is not taking effect. PR1597848
- Socket connection drops due to keepalive timer expiration with port 33015.PR1598019
- s-Flow impacts ICMP traffic on the QFX5120-48Y switches. PR1598239
- On QFX5100, QFX5110, QFX5120, QFX5200, and QFX5210 platforms, DDoS violations might be reported incorrectly for IP multicast miss traffic (IPMCAST-MISS). PR1598678
- File permissions are changed for /var/db/scripts files after reboot. PR1599365
- The Layer 3 traffic gets silently dropped and discarded on the QFX10002-60C devices with IRB interface.PR1599692
- The VCP might not form adjacency after rebooting the primary FPC in Virtual Chassis scenario. PR1600398
- Unable to disable the management port em1. PR1600905
- Removing and adding Virtual Chassis ports might cause the FPC to reboot. PR1601557
- On QFX5120-48Y switches, dc-pfe core file is generated while issuing the `show pfe vxlan nh-usage` command in an ERB EMC scenario with ~6000 ARP entries.PR1601949
- The IPv6 traffic might be impacted on the QFX platforms when an IPv6 route resolves over a dynamic tunnel. PR1602007
- In rare cases, l2ald might abort when EVPN/VxLAN interface related configuration is deleted. PR1602244
- The egress interface of the GRE tunnel is not dynamically updated when the destination to tunnel changes. PR1602391

- In Junos OS, specific packets over VXLAN cause FPC memory leak and ultimately reset (CVE-2022-22170). PR1602407
- The status of FPC becomes down and the dcpfe process might generate core file dump in some cases. PR1602583
- In Junos OS, an l2cpd memory leak might occur when specific LLDP packets are received leading to a DoS (CVE-2022-22172). PR1602588
- Traffic loss might be seen in MC-LAG scenario. PR1602811
- On the QFX5000 line of switches, traffic drop might occur in the Virtual Chassis scenario when you configure the firewall filter. PR1602914
- The dot1x authentication may not work on EVPN/xlan enabled endpoints. PR1603015
- Packet loss might be seen on filter-based GRE deployments. PR1603453
- In Junos OS, QFX5000 Series devices might run out of memory, causing traffic loss, upon receipt of specific IPv6 packets (CVE-2022-22174). PR1603531
- The l2ald might crash when ESI with local interface goes down in EVPN-VXLAN scenario. PR1603979
- Duplicate packets might be seen during bringing up all the interfaces on the spine switch. PR1604393
- The carrier transition counter might not get incremented upon link flap after the reboot. PR1605037
- MAC move might be seen between the ICL and the MC-LAG interface if you add or remove VLANs on the ICL interface. PR1605234
- On QFX5100, generate an optical power after detached and attached QSFP on disabled interface. PR1606003
- Multicast streams might stop flooding in VXLAN setup. PR1606256
- The VCP ports goes down after the members split and merge in the Virtual Chassis. PR1606705
- The LLDP packets received on the VXLAN enabled port might get flooded unexpectedly. PR1607249
- On QFX Series switches, the fxpc process might crash and generate a core dump. PR1607372
- FPC might crash post firewall filter configuration changes in QFX platforms. PR1608610
- An additional VLAN tag might be added for PPPoE (Point-to-Point Protocol over Ethernet) packets on QFX10016. PR1610012
- Ping to lo0/IRB over Type-5 fails. PR1610093

- On QFX10002 line of switches, continuous Layer 3 traffic drop occurs with the MC-LAG configuration.PR1610173
- On QFX5000 Virtual Chassis, MAC move or MAC flap might be triggered. PR1610295
- Inter-vlan connectivity might be lost in an EVPN-VXLAN with CRB topology. PR1611488
- Layer 3 interfaces unable to attach DSCP rewrite firewall filter on QFX5100. PR1612587
- On QFX10002-60C line of switches, continuos FPC might crash and the dcpfe process might generate the core file. PR1612871
- Arp resolution for data traffic received over Type5 might fail. PR1612905
- FPC might crash after device restart in EVPN-VXLAN scenario. PR1613702
- Removing the optical module "JNP-SFPP-10GE-T" from a port might cause certain ports to go down. PR1614139
- On QFX5000 VLAN firewall filter is not deleted in Packet Forwarding Engine after configuration change. PR1614767
- The l2ald process might crash in EVPN scenario. PR1615269
- Slow memory leak (32 bytes each time) of rpd might be seen. PR1616065
- Packet drop might occur if VxLAN is configured. PR1616683
- BGP routes learnt through type 5 EVPN routes might not get activated. PR1617878
- BFD session might get stuck in initial state after l2-learning restart due to incomplete ARP resolutions. PR1618280
- Core dumps might be seen on QFX devices after configuration changes. PR1618352
- Dot1x based firewall policers are not supported. PR1619405
- The process dcpfe might crash after performing VXLAN VNI configuration change and delete on QFX5000 platforms. PR1619445
- Disabled VCP (Virtual chassis port) will be up after the optic on it is reseated. PR1619997
- High wired memory utilization might be observed if GRES is enabled. PR1620599
- Routes learned through the EVPN Type-5 route are not resolved. PR1620627
- EVPN-VXLAN Type5 traffic might get failed on the Spine device of QFX10000. PR1620924
- Implement `show task scheduler-slip-history` to display no of scheduler slips and last 64 slip details. PR1626148

- JDI-RCT: QFX10002 MCLAG PDT: L3 Traffic failures observed continuously with PDT mclag configuration. PR1627846
- 802.1p BA classification might not work on mixed Virtual Chassis when interface has a DSCP and 802.1p classifier. PR1628447
- The vmhost crash might be seen in a rare condition when route addition and change. PR1629200
- Data might not be exchanged through EVPN-VxLAN domain. PR1635347
- Memory usage continuously increase is observed on backup chassis if subscriber service is enabled. PR1595238

Class of Service (CoS)

- Transit packets from local to remote VTEP might get punted to CPU and cause DDoS events. PR1489233
- The dcpfe crash might be seen on QFX5120. PR1563625
- The TCP-ECN traffic might not be forwarded with high priority. PR1585854
- The dcpfe core might be seen in auto-channelization scenario or when SFP is plugged out. PR1616847

EVPN

- Traffic loss might be seen under EVPN scenario when MAC-IP moves from one CE interface to another. PR1591264
- The device announces router-mac, target, and EVPN VXLAN community to BGP IPv4 NLRI. PR1600653
- Traffic sent by the QFX5000 switch leaf to remote leaf with link down. PR1605375

High Availability (HA) and Resiliency

- During ISSU package signature validation might fail and the upgrade might not happen. PR1575680
- Memory leaking might occur on backup Routing Engine when ksyncd is in inconsistent state and had encountered an initialization error. PR1601960

Layer 2 Ethernet Services

- The DHCP client might be offline for about 120 seconds after sending the DHCPINFORM message. PR1587982

MPLS

- Traffic loss seen on QFX5000 after STP topology change. PR1616878

Platform and Infrastructure

- IPv6 link-local traffic is getting classified to firewall host this might affect communication on IPv6 link-local addresses. PR1600085

Routing Policy and Firewall Filters

- The interface-routes rib-group policy does not work as expected in the VxLAN scenario. PR1537306
- The rpd process might get stuck at 100% when EVPN vrf-target is enabled and after any configuration change. PR1616167

Routing Protocols

- The remaining BFD sessions of the aggregated Ethernet interface flap continuously if one of the BFD sessions is deleted. PR1516556
- IPv4 static route might still forward traffic unexpectedly even when the static route configuration has already been deleted. PR1599084
- The interface might receive multicast traffic from a multicast group which it is not interested in. PR1612279
- The wrong BGP path might get selected even when a better or preferred route is available. PR1616595
- Time delay to export prefixes to BGP neighbors might occur post applying peer-specific BGP export policies. PR1626367

Resolved Issues: 21.1R2

IN THIS SECTION

General Routing

- Kernel crash might occur after NSSU while performing GRES. PR1533874
- The dcpfe process might crash and cause FPC to restart due to the traffic burst. PR1534340
- FPC(s) might not boot-up on MX960 and EX9214 in a certain condition. PR1545838
- The dcpfe process might crash on QFX10000 platforms. PR1546572
- The traffic will not be load-balanced properly in EVPN overlay-ecmp setup. PR1550020
- The interface might not come up with 1G optics. PR1554098
- The dcpfe process might crash and restart with a dcpfe core file created while running the Type5 EVPN-VXLAN with 2000 VLANs. PR1556561
- In QFX platforms, Layer 3 static license is required though it is included in base license. PR1557631
- The MAC addresses learned in a Virtual Chassis might fail due to aging out in the MAC scaling environment. PR1558128
- The VCF might become not stable. PR1559172

- The subscriber management infrastructure daemon (smid) process might be stuck at 100%. PR1559402
- On the QFX5110 line of switches, the untagged traffic routed over native-vlan might be dropped. PR1560038
- On the QFX5200 line of switches, the pseudorandom binary sequence (PRBS) test fails for 100GbE interfaces with the default settings. PR1560086
- When configuring the static MAC and static ARP on the EVPN core aggregate interface, the underlay next-hop programming might not be updated in the Packet Forwarding Engine. PR1561084
- The tunable optics SFP+-10G-T-DWDM-ZR does not work. PR1561181
- Dcpfe process might crash on after committing EVPN-VXLAN profile configuration and ARP resolution might fail causing traffic issues. PR1561588
- Junos OS, QFX5000 Series platform, traffic from the network internal to the device (128.0.0.0) might be forwarded to egress interfaces. (CVE-2021-31371) PR1561722
- On QFX5000 platforms, the dcpfe process might crash after deleting VXLAN configuration. PR1562692
- QFX5110-48s-4c :: ptp traffic-statistics are not as expected. PR1563876
- On the QFX5100 Virtual Chassis, the following continuous message is observed: agentd-pfe-proxy_telemetry_publisher. PR1566528
- On the QFX5100 line of switches, the following internal comment is displayed: Placeholder for QFX platform configuration. PR1567037
- RPD core dump is observed at device reboot and/or daemon restart time. Daemon recovers and there is no service impact on routing protocol usage. PR1567043
- On the QFX10002 line of switches, discrepancy in inet.1 versus Packet Forwarding Engine reports multicast routes. PR1567353
- MAC addresses might not be relearned successfully after MAC address age timeout. PR1567723
- The 100G port with module QSFP 100G-SR4-T2 converts to two channelized interfaces without any channelized configuration. PR1567937
- Another port will also be shutdown after shutting down one port on QFX10002-60C. PR1568294
- BFD flaps seen between leaf and core during spine reboot causing other protocols flap. PR1568615
- On the QFX10000 line of switches, the firewall log is incorrectly populating from the Packet Forwarding Engine for IPv6 traffic. PR1569120

- The dcpfe might crash if the TYPE-5 tunnel is failed to be installed for EVPN-VxLAN. PR1570136
- Junos OS, QFX5100-96s platform, major alarm set, fan x does not spin. The fan logs are seen. PR1570587
- PTP management message with SMTLV is sent only to the first port number to go active in the member multicast-mode l2-ifl. PR1571283
- Unexpected packet loss might occur if you delete the subunit of the physical interface. PR1571286
- The dcpfe crash is seen after running MC-LAG profile configuration. PR1571471
- DCI traffic loss of 100% observed in transit spine devices. PR1572238
- EVPN VXLAN CE interface with RSTP configured might cause LACP or BFD issues. PR1572504
- On the QFX10008 chassis, the dcpfe process generates a core file. PR1572889
- Traffic loss might be observed due to faulty FPC on QFX10008 and QFX10016 platforms. PR1574779
- Port-mirroring might not work when the analyzer output is a trunk interface. PR1575129
- On the QFX10000 line of switches, a high rate of 802.3X pause frames are sent out of the interfaces. PR1575280
- BFD flaps might be seen occasionally during spine reboot. PR1575296
- On QFX Series switches, upgrading to version 20.3 or later might report a **warning: requires 'l3vpn' license** message on commit when a VRF instance configuration exists. PR1575608
- The dual-speed supported DAC cable (100G to 4x25G Splitter) might not come up on QFX5120-48Y. PR1576180
- Analyzer is not working on all Junos QFX5000 platforms. PR1576327
- On Junos OS QFX5000 Series, control traffic might be dropped if a high rate of specific multicast traffic is received (CVE-2021-31370). PR1576488
- The OSPF session over IRB might not come up in the EVPN-VXLAN scenario. PR1577183
- The WAN port links might not get brought down immediately during some abnormal type of linecard reboot on QFX10000 platforms. PR1577315
- TACACS traffic might be dropped. PR1578579
- The ISIS packet might be corrupted on the provider edge device over the layer 2 circuit tunnel. PR1580047

- The dcpfe process crashes while checking the virtual tunnel-nh packet status. PR1580114
- DHCP packets might be dropped if dynamic filter **dyn-dhcpv4_v6_trap** is applied on the interface. PR1580352
- While mapping analyzers to the channelized port, mirror might not work properly. PR1580473
- On the QFX5120-32C line of switches, the following error is observed: **kern.ipc.maxpipekva exceeded; see tuning error**. PR1581192
- The switchover might be affected with the shared VXLAN tunnel. PR1581524
- The traffic might not be load-balanced properly in an EVPN overlay-ecmp setup. PR1582017
- Some 40G ports might not be channelized successfully on the QFX5100 platforms. PR1582105
- On the QFX10000 line of switches, the firewall filter logs are incorrectly populated the protocol 8847 entries. PR1582780
- Firewall filter not programmed after deleting a large filter and adding a new one in a single commit on QFX5000 platforms. PR1583440
- The ZTP process might cause the black holing of the traffic. PR1585057
- DHCP offer packets might be dropped on Spine device in VxLAN multi-homing setup. PR1585715
- Inter and intra VNI traffic drop might occur in spine with EVPN-VxLAN CRB configuration. PR1586537
- On Junos OS, an FPC heap memory leak will be triggered by certain Flowspec route operations which can lead to an FPC crash (CVE-2021-31367). PR1589133
- 50% traffic loss might happen in EVPN-VxLAN scenario. PR1589547
- When the deleted aggregated Ethernet interface is not getting deleted (mirror trunk group) in the hardware for the analyzer input aggregated Ethernet interface. PR1589579
- LLDP packets drop on SP style interface for QFX devices. PR1589702
- VXLAN DDoS violation might occur when disabling the port mirror analyzer output interface. PR1590150
- VC primaryship changed and connection dropped after renumbering of backup member ID. PR1590358
- On QFX5120-48T, after removing 1g speed on interfaces it won't come back as 10g. PR1591038
- xSTP might not get configured when enabled on a interface with SP style configuration on all platforms. PR1592264

- Routing Engine kernel might crash due to IFL of aggregated interface adding failure in Junos kernel. PR1592456
- The IPv4 fragmented packets might be broken if PTP transparent clock is configured. PR1592463
- MPLS traffic might get discarded on passive monitoring interface on QFX10002, QFX10008 and QFX10016 switches. PR1592693
- Multiple crashes with **toe_interrupt_errors** might be observed. PR1593025
- BFD session might flap during Routing Engine switchover. PR1593244
- The dcpfe process might crash in EVPN-VxLAN scenario. PR1593950
- Packet drop might occur in ECMP next-hop flap scenario. PR1594030
- The existing ECMP route traffic might be dropped if configuring a static ECMP route with the same number of next-hops as the existing ECMP route. PR1594573
- The re-installation of the type-5 tunnels might fail in the EVPN-VXLAN scenario. PR1595197
- The **fpc0 bcm pkt reinsert failed** log written in the log messages in an aggressive way. PR1596643
- Deletion of macsec configuration on an IFL is not taking effect. PR1597848
- Sflow impacts on ICMP traffic on QFX5000 series platforms. PR1598239
- The layer 3 traffic blackholing might be seen on QFX10002-60C devices with IRB interface. PR1599692
- FPC down and dcpfe core dump migtht be seen in some cases. PR1602583
- Traffic loss might be seen in MC-LAG scenario on QFX platforms. PR1602811
- The dot1x authentication may not work on EVPN/xlan enabled endpoints. PR1603015
- Duplicate packets might be seen during bringing up all the interfaces on the spine switch. PR1604393
- 20.2R2-S2 QFX5000 MAC moves between ISL and MC-AE LAGs and flooding after upgrade and adding VLANs. PR1605234
- Multicast streams might stop flooding in VXLAN setup. PR1606256

Class of Service (CoS)

- Traffic might be dropped by destination device. PR1568333
- Unable to configure policer with bandwidth-limit greater than 50g. PR1575049

- The buffer allocation for VCP ports might not get released in Packet Forwarding Engine after physically moving the port location. PR1581187
- The dscp classifier does not work and all packets are sent to a single queue. PR1585361

EVPN

- global-mac-ip-table-aging-time changes from a high to low value might not take effect. PR1562925
- The dev-longevity l2ald process generates the core file at l2ald_next_bd_member. PR1570757
- Configuring static-mac and no-mac-learning simultaneously on the VXLAN interface causes stale MAC/IP entry in the EVPN database. PR1576147
- After device reboot in EVPN-VXLAN setup with graceful restart, EVPN routes are not advertised to EVPN peers until rpd is up for 180 seconds. PR1586246

Interfaces and Chassis

- MAC address entry issue might be observed after the MC-LAG interface. PR1562535
- New added MC-LAGs do not come up after Routing Engine switchover. PR1583547
- Removing the configuration from interface stanza might cause the dcpfe process to crash. PR1594356

Layer 2 Features

- On the QFX5110-32Q switches, LACP does not come up in the non-oversubscribed mode for a set of ports. PR1563171
- Traffic forwarding for VLAN 2 might not be correct when a VLAN member is removed from the ESI interface. PR1570446
- The dcpfe crashes in the VxLAN scenario. PR1571170
- The DF might not forward BUM traffic on QFX5000 series switches. PR1575976
- MAC addresses learnt from the MC-LAG client device might keep flapping between the ICL interface and MC-AE interface after one child link in the MC-AE interface is disabled. PR1582473
- Traffic drop might be seen on the aggregated Ethernet interface. PR1585320

Layer 2 Ethernet Services

- DHCP packet drop might be observed when the DHCP relay is configured on a leaf device. PR1554992
- The DHCP client will be offline for 120 seconds after sending the DHCPINFORM message in the DHCP relay scenario. PR1575740
- DHCP relay drops packets during the renewal DHCP process. PR1576417

Network Management and Monitoring

- Slow memory leak could be observed for snmpd process. PR1575790

Platform and Infrastructure

- Ex3400 VC - Console access on backup VC member is not allowed. PR1530106
- Junos OS: Upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart (vmcore) (CVE-2021-0283, CVE-2021-0284). PR1557881
- FPC might crash in a scaled-firewall configuration. PR1586817

Routing Policy and Firewall Filters

- The rpd might crash when the deletion of routing table occurs. PR1565629

Routing Protocols

- Traffic might be silently discarded when a BGP route that is part of multipath gets deleted. PR1514966
- The BFD sessions over IRB interface stuck in Init state with FRR errors incrementing. PR1541851
- The fxpc process might crash after flapping the related protocols in the ECMP scenario. PR1556224
- The ppmd memory leak might cause traffic loss. PR1561850
- There might be traffic loss when GRE interface flaps on QFX platforms. PR1566428
- On QFX5000 platforms memory leak might be observed. PR1566483
- The untagged packets might not work on EX Series platforms. PR1568533
- Memory leak might happen in MSDP scenario. PR1571906

- The GRE egress traffic might not be forwarded between the different routing instances. PR1573411
- The QFX5000 line of switches might drop the DHCP packets in the static VXLAN scenario. PR1576168
- Multicast Packets with TTL=1 are dropped on VXLAN enabled interface when igmp-snooping/MLD-snooping is enabled. PR1576775
- Traffic loss might be observed in the EVPN-VxLAN scenario on QFX5000 platforms. PR1580005
- BGP session carrying VPNv4 prefix with IPv6 next-hop might be dropped. PR1580578
- The dcpfe process might crash when any interface flaps. PR1579736
- Traffic loss might be seen when ipv6 traffic forwarded by ipv4 GRE tunnel. PR1582408
- With IGMP snooping implemented, there is unexpected jitter issue that could cause traffic loss. PR1583207
- The rpd process might crash after committing with the configured static group 224.0.0.0. PR1586631
- BGP Egress-TE routes lose to BGP routes using the same protocol-preference. PR1593332

Resolved Issues: 21.1R1

EVPN

- All the ARP reply packets toward some address are flooded across the entire fabric. PR1535515
- EVPN-VXLAN registers MAC-move counters under system statistics bridge even though there is no actual MAC move for the multihomed clients. PR1538117
- The l2ald process might generate a core file when changing the EVPN-VXLAN configuration. PR1541904
- The l2ald deamon might crash when `forwarding-options evpn-vxlan shared-tunnels` is configured. PR1548502
- The l2ald process generates a core file at l2ald_iff_rtm_delete_subintf_ifbds during dci fusion run. PR1550109
- QFX10002 :: mac-vrf: QFX10k l2ald core: l2ald_vxlan_ifl_create_event_handler at /src/junos/usr.sbin/l2ald/platform/junos/l2ald_rtsock_vxlan.c:477 PR1560068
- [evpn_vxlan]: evpn vxlan mac-ip aging testcase failed. PR1562925

- l2ald process generates a core file at vlogging_event, l2ald_vxlan_ifl_create_event_handler, l2ald_vxlan_ifl_event_handler, l2ald_process_event. PR1576558

Forwarding and Sampling

- The l2ald process might crash due to a next-hop issue in the EVPN-MPLS. PR1548124
- Configuration archive transfer-on-commit fails on Junos OS Release 18.2R3-S6.5. PR1563641

General Routing

- Port qualifier is not supported for QFX5000 platforms. PR1440980
- On the QFX5000 line of switches, the egress ACL filter entries is only 512 in Junos OS Release 19.4R1. PR1472206
- On the QFX10000 device, the chassisd process might generate core files on the backup Routing Engine after commit for 200 seconds due to the following error message: CHASSISD_MAIN_THREAD_STALLED. PR1481143
- On the QFX5000 line of switches, multicast traffic loss is observed due to a few missing multicast routes in the spine node. PR1510794
- The DHCP traffic might not be forwarded correctly when DHCP sends unicast packets. PR1512175
- Channelized interfaces might fail to come up. PR1512203
- The output of the show chassis forwarding-options command displays incorrect display issue, Virtual Chassis environment, and configured num-65-127-prefix values. PR1512712
- On the QFX5100 device, the cprod process timeout triggers high CPU utilization. PR1520956
- Output interface index in the sFlow packet is zero when transit traffic is observed on the IRB interface with VRRP enabled. PR1521732
- Some inter-VLAN traffic flows do not converge after rebooting a spine (QFX10002) device in an EVPN-VXLAN non-collapsed scaled scenario. PR1522585
- Traffic loss might be observed on interfaces in a VXLAN environment. PR1524955
- Channelizing the 40GbE port to a 10GbE port might bring down another interface on the QFX10000 platforms. PR1527814
- When a multicast feed is received with TTL 1 on QFX10002 line of switches. There will be 2 copies of the packet sent to the host - one from the normal flow and another from the multicast module. These packets are logged in the firewall log incorrectly. PR1533814

- The dcpfe process might crash and cause FPC to restart due to the traffic burst. PR1534340
- High rate of ARP or NS packets might be observed between a device that runs Junos OS and the host when the device that runs Junos OS receives an ARP or NS packet on an interface in transition. PR1534796
- The following Packet Forwarding Engine error message is observed in the BRCM-VIRTUAL: brcm_virtual_tunnel_port_create() ,489: Failed NW vxlan port token(45) hw-id(7026) status(Entry not found). PR1535555
- The interfaces on QFX5100-48T switch might stay up when the peer device is rebooting PR1538071
- On the QFX5100-48T, interfaces are not created after a channel speed of 10 Gbps is applied on ports 48 through 53. PR1538340
- The BFD sessions might not come up in an VXLAN scenario. PR1538600
- Management Ethernet link down alarm is seen while verifying system alarms in a Virtual Chassis setup. PR1538674
- ARP request may be dropped in leaf node in an EVPN-VXLAN scenario. PR1539278
- The rpd memory leak might be observed on the backup Routing Engine due to link flaps PR1539601
- Unable to take RSI properly due to the authentication error. PR1539654
- FPC might not be recognized after power cycle (hard reboot) PR1540107
- Traffic loss might be seen in the OVSDB VXLAN scenario. PR1540208
- The Packet Forwarding Engine might crash in MPLS IPv6-tunneling scenario when the next hop changes. PR1540793
- On the QFX5100 Virtual Chassis, the End Segment Not Present message is not reported for the ping overlay function with the local host MAC. PR1542226
- On the QFX5000 device running EVPN-VXLAN, the following Packet Forwarding Engine error message might be seen: bd_platform_irb_ifl_attach_detach: platform specific irb ifl attach/detach failed (-1). PR1543812
- On the QFX10002-60C device, the show pfe filter command is unavailable. PR1545019
- The chip on FPC linecard might crash during the system booting. PR1545455
- OSPFv3 session may keep flapping and OSPFv3 hellos might be dropped in the host path. PR1547032

- On a QFX10000 device, traffic might get dropped when the set routing-options forwarding-table no-ecmp-fast-reroute configuration is changed to 128 ECMP entries. PR1547457
- On the QFX5100 Virtual Chassis, the backup Routing Engines clear the reporting alarm for a PEM failure intermittently for a missing power source. PR1548079
- The VXLAN encapsulated packet might be sent on the network port with an incorrect inner VLAN ID 4095. PR1548218
- The 40GbE interface might be channelized after restarting the Virtual Chassis member. PR1548267
- The Neighbor Solicitation might be dropped from the peer device. PR1550632
- The interface filter with source-port 0 matches everything instead of port 0. PR1551305
- On the QFX5110 and QFX5120 devices, the DHCPv6 traffic received over the VTEP might not be forwarded. PR1551710
- On the QFX5000 devices, ARP resolution might fail. PR1552671
- The `action-shutdown` configuration of storm control does not work for ARP broadcast packets. PR1552815
- Traffic might not passed due to the addition of the VLAN tag 2 while passing through the Virtual Chassis port. PR1555835
- Traffic might be dropped when a firewall filter rule uses the then VLAN action. PR1556198
- The dcpfe process crashes and a core file is generated on QFX10002-60C while running Type 5 EVPN_VXLAN configuration with 2000 VLANs. PR1556561
- DHCP Discover packets are not getting flooded with VXLAN configuration. PR1557049
- Traffic storm might be caused by the analyzer due to link flapping. PR1557274
- Firewall filter might fail to work on QFX5000 platforms. PR1558320
- On the QFX5120 device, amber LEDs are displayed for the fan modules after upgrading to Junos OS Release 20.2R1. PR1558407
- Pseudo Random Binary Sequence (PRBS) test on QFX5200 platform fails for 100GbE interfaces with default settings. PR1560086
- There are a few instances of IPv6 ARP ND failure after loading the base configurations. PR1560161
- When configuring static MAC and static ARP on the EVPN core aggregate interface, the underlay next-hop programming might not be updated in the Packet Forwarding Engine. PR1561084
- PTP boundary clock with G.8275.2.enh profile_2 512 clients does not come up. PR1561348

- PTP lock status gets stuck at the Acquiring state instead of the Phase Aligned state. PR1561372
- Firewall filters might not be working after ISSU. PR1561690
- On QFX10000 platforms, the dcpfe process might crash during configuration changes. PR1561746
- Traffic loss might happen in a large-scaled EVPN scenario when the next-hop type changes between Discard and Unicast. PR1562425
- Port mirroring might not work as expected on QFX5000 platforms. PR1562607
- Due to a failure in an FPGA image load, the PTP BC/OC application fails to send and receive packets properly, resulting in full failure of the PTP BC/OC operation. PR1563876
- Output of the `show chassis fpc ether-types` command includes the FPC slot number. PR1564496
- QFX10K: Firewall log incorrectly populating from PFE for IPv6 traffic. PR1569120
- QFX10002:OpenConfig to Junos OS configuration translation has failed while translating interface-mode trunk and vlan members. PR1580292

Interfaces and Chassis

- The logical interface might flap after the addition or deletion of the native VLAN configuration. PR1539991
- MAC address entry issue might be seen after MC-LAG interface failover or failback. PR1562535

Layer 2 Ethernet Services

- DHCP packet drop may be seen when DHCP relay is configured on the leaf device. PR1554992

Layer 2 Features

- Check traffic with VXLAN encapsulation header fails.PR1541316
- Traffic may be forwarded incorrectly on an interface having VXLAN enabled and "hold-time up xxx" statement configured. PR1550918
- On QFX5000 and EX4600 platforms, memory leak might happen when the physical interface is continuously attaching and detaching. The dcpfe might crash if the device is running out of memory. Traffic loss might be seen during the dcpfe crash and restart. PR1543169
- On EX4650-48Y and QFX5120 platforms, packets with VLAN ID 0 are dropped. PR1566850

Routing Policy and Firewall Filters

- The policy configuration might be mismatched between the rpd and mgd process when deactivate policy-options prefix-list is involved in the configuration sequence. PR1523891

Routing Protocols

- On the QFX5100-48T-6Q Virtual Chassis or Virtual Chassis fabric, the following error message is observed while copying the image to the Virtual Chassis fabric member and trying to downgrade the image: rcp for member 14, failed. PR1486632
- The IPv6 traffic might be silently dropped due to null-route filtering when falling back from IP-in-IP tunnel to inet.0/inet6.0. PR1508631
- Traffic might be silently discarded when the clear bgp neighbor all command is executed on a router and also on the corresponding route reflector in succession. PR1514966
- The OSPF neighborship gets stuck in the Start state after configuring the EVPN-VXLAN. PR1519244
- DCPFE crash might be observed while updating VRF for multicast routes during IRB uninit. PR1546745
- [pfe] [generic] : qfx5100-24q-2p :: fxpc core in brcm_nh_unilist.c:2162 during stress test PR1556224
- BGP-LU session flap might be seen with AIGP used scenario. PR1558102
- On the QFX5110-32Q device, the following syslog error message is observed after loading the NC T5 EVPN-VXLAN configuration: BCM-L2,pfe_bcm_l2_sp_bridge_port_tpid_set() Config TPID New/Old (8100:8100) Other-Tpid's ba49, 4aa0, 80f. PR1558189
- Layer 3 inter pod IPv4 traffic issue observed after loading non-collapsed Type 5 EVPN-VXLN configuration. PR1560173
- On the QFX5110 platform, ARP resolution may fail if "native-vlan-id" is configured on the VXLAN interface. PR1563569
- The dcpfe process might crash when the size of the Local Bias Filter Bitmap string exceeds 256 characters. PR1568159
- On the QFX5210-64C device, ping does not work while verifying the native VLAN behavior on the Q-in-Q interface. PR1568533

User Interface and Configuration

- set chassis fpc 0 ether-type applicable only for ether index 6 to 27. PR1565695

Virtual Chassis

- On the QFX5000 Virtual Chassis, the DDoS violations that occur on the backup are not reported to the Routing Engine. PR1490552

Documentation Updates

There are no errata or changes in Junos OS Release 21.3R1 documentation for QFX Series documentation.

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 **20.3** in the Release pull-down list to the right of the Software tab on the Download Software page.

4. In the Install Package section of the Software tab, select the QFX Series Install Package for the 20.3 release.

   An Alert box appears.

5. In the Alert box, click the link to the PSN document for details about the software, and click the link to download it.

   A login screen appears.

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

7. Download the software to a local host.

8. Copy the software to the device or to your internal software distribution site.

9. Install the new jinstall package on the device.

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

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

```
user@host> request system software add source/jinstall-host-qfx-5-x86-64-20.3-R1.n-secure-signed.tgz reboot
```

Replace *source* with one of the following values:

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

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

Rebooting occurs only if the upgrade is successful.

NOTE: After you install a Junos OS Release 20.3 `jinstall` package, you can issue the `request system software rollback` command to return to the previously installed software.

Installing the Software on QFX10002-60C Switches

This section explains how to upgrade the software, which includes both the host OS and the Junos OS. This upgrade requires that you use a VM host package—for example, a **junos-vmhost-install-*x*.tgz** .

During a software upgrade, the alternate partition of the SSD is upgraded, which will become primary partition after a reboot .If there is a boot failure on the primary SSD, the switch can boot using the snapshot available on the alternate SSD.

NOTE: The QFX10002-60C switch supports only the 64-bit version of Junos OS.

NOTE: If you have important files in directories other than /config and /var, copy the files to a secure location before upgrading. The files under /config and /var (except /var/etc) are preserved after the upgrade.

To upgrade the software, you can use the following methods:

If the installation package resides locally on the switch, execute the **request vmhost software add** ***<pathname><source>*** command.

For example:

```
user@switch> request vmhost software add /var/tmp/junos-vmhost-install-qfx-x86-64-20.4R1.9.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-x86-64-20.4R1.9.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 18.3R1.

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-x86-64-20.4R1.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-x86-64-20.4R1.n-secure-signed.tgz reboot
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the `show version` command.

```
user@switch> show version
```

Upgrading Software from Junos OS Release 15.1X53-D3X to Junos OS Release 15.1X53-D60, 15.1X53-D61.7, 15.1X53-D62, and 15.1X53-D63 on QFX10008 and QFX10016 Switches

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

The switch contains two Routing Engines, so you will need to install the software on each Routing Engine (re0 and re1).

If the installation package resides locally on the switch, execute the **request system software add** ***<pathname><source>*** command.

To install the software on re0:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

If the Install Package resides remotely from the switch, execute the **request system software add** ***<pathname><source>*** **re0** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re0
```

To install the software on re1:

```
user@switch> request system software add /var/tmp/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

If the Install Package resides remotely from the switch, execute the **request system software add** ***<pathname><source>*** **re1** command.

For example:

```
user@switch> request system software add ftp://ftpserver/directory/jinstall-host-qfx-10-m-15.1X53-D60.n-secure-domestic-signed.tgz re1
```

Reboot both Routing Engines.

For example:

```
user@switch> request system reboot both-routing-engines
```

After the reboot has finished, verify that the new version of software has been properly installed by executing the `show version` command.

```
user@switch> show version
```

Installing the Software on QFX10008 and QFX10016 Switches

Because the switch has two Routing Engines, perform a Junos OS installation on each Routing Engine separately to avoid disrupting network operation.

NOTE: Before you install the software, back up any critical files in **/var/home**. For more information regarding how to back up critical files, contact Customer Support at https://www.juniper.net/support.

WARNING: If graceful Routing Engine switchover (GRES), nonstop bridging (NSB), or nonstop active routing (NSR) is enabled when you initiate a software installation, the software does not install properly. Make sure you issue the CLI `delete chassis redundancy` command when prompted. If GRES is enabled, it will be removed with the `redundancy` command. By default, NSR is disabled. If NSR is enabled, remove the nonstop-routing statement from the `[edit routing-options]` hierarchy level to disable it.

1. Log in to the master Routing Engine's console.

   For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

2. From the command line, enter configuration mode:

   ```
   user@switch> configure
   ```

3. Disable Routing Engine redundancy:

   ```
   user@switch# delete chassis redundancy
   ```

4. Disable nonstop-bridging:

```
user@switch# delete protocols layer2-control nonstop-bridging
```

5. Save the configuration change on both Routing Engines:

```
user@switch# commit synchronize
```

6. Exit the CLI configuration mode:

```
user@switch# exit
```

After the switch has been prepared, you first install the new Junos OS release on the backup Routing Engine, while keeping the currently running software version on the master Routing Engine. This enables the master Routing Engine to continue operations, minimizing disruption to your network.

After making sure that the new software version is running correctly on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the software version on the other Routing Engine.

7. Log in to the console port on the other Routing Engine (currently the backup).

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

8. Install the new software package using the `request system software add` command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-
x86-64-20.4R1.n-secure-signed.tgz
```

For more information about the `request system software add` command, see the CLI Explorer.

9. Reboot the switch to start the new software using the `request system reboot` command:

```
user@switch> request system reboot
```

NOTE: You must reboot the switch to load the new installation of Junos OS on the switch.

To abort the installation, do not reboot your switch. Instead, finish the installation and then issue the `request system software delete <package-name>` command. This is your last chance to stop the installation.

All the software is loaded when you reboot the switch. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation is not sending traffic.

10. Log in and issue the `show version` command to verify the version of the software installed.

```
user@switch> show version
```

Once the software is installed on the backup Routing Engine, you are ready to switch routing control to the backup Routing Engine, and then upgrade or downgrade the master Routing Engine software.

11. Log in to the master Routing Engine console port.

For more information about logging in to the Routing Engine through the console port, see the specific hardware guide for your switch.

12. Transfer routing control to the backup Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the `request chassis routing-engine master` command, see the CLI Explorer.

13. Verify that the backup Routing Engine (slot 1) is the master Routing Engine:

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
    Current state                  Backup
    Election priority              Master (default)

Routing Engine status:
  Slot 1:
```

```
    Current state                   Master
    Election priority               Backup (default)
```

14. Install the new software package using the `request system software add` command:

```
user@switch> request system software add validate /var/tmp/jinstall-host-qfx-10-f-
x86-64-20.4R1.n-secure-signed.tgz
```

For more information about the `request system software add` command, see the CLI Explorer.

15. Reboot the Routing Engine using the `request system reboot` command:

```
user@switch> request system reboot
```

NOTE: You must reboot to load the new installation of Junos OS on the switch.
To abort the installation, do not reboot your system. Instead, finish the installation and then issue the `request system software delete jinstall` *<package-name>* command. This is your last chance to stop the installation.

The software is loaded when you reboot the system. Installation can take between 5 and 10 minutes. The switch then reboots from the boot device on which the software was just installed. When the reboot is complete, the switch displays the login prompt.

While the software is being upgraded, the Routing Engine on which you are performing the installation does not send traffic.

16. Log in and issue the `show version` command to verify the version of the software installed.

17. Transfer routing control back to the master Routing Engine:

```
user@switch> request chassis routing-engine master switch
```

For more information about the `request chassis routing-engine master` command, see the CLI Explorer.

18. Verify that the master Routing Engine (slot 0) is indeed the master Routing Engine:

```
user@switch> show chassis routing-engine
Routing Engine status:
  Slot 0:
```

```
    Current state                  Master
    Election priority              Master (default)

outing Engine status:
  Slot 1:
    Current state                  Backup
    Election priority              Backup (default)
```

Performing a Unified ISSU

You can use unified ISSU to upgrade the software running on the switch with minimal traffic disruption during the upgrade.

NOTE: Unified ISSU is supported in Junos OS Release 13.2X51-D15 and later.

Perform the following tasks:

- No Link Title
- No Link Title

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-x86-64-20.4R1.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-x86-64-18.3R1.n-secure-signed.tgz ...
Install jinstall-host-qfx-5-f-x86-64-19.2R1.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 and Extended End-Of-Life Releases

We have two types of releases, EOL and EEOL:

- End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.

  Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 17: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Junos OS Release Notes for SRX Series

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for the SRX Series Services Gateways. 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for SRX Series devices.

What's New in 21.1R3

There are no new features for SRX Series Services Gateways in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features for SRX Series Services Gateways in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for SRX Series devices.

Application Identification (AppID)

- **Application signature package enhancements (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 21.1R1, we've enhanced the application signature package by grouping all newly added signatures under the junos:all-new-apps group. When you download the application signature package on your device, the predefined application group is downloaded. You can use this application group in the security policy configuration.

  We've also introduced a list of application tags, based on attributes, in the application signature package. You can group similar applications based on these predefined tags. By doing so, you can consistently reuse the application groups when you define security policies.

  [See Predefined Application Signatures for Application Identification.]

- **Enhancements to packet capture of unknown applications (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 21.1R1, your security device stores the packet capture of unknown

  applications' details per session. As a result of this change, the packet capture (.pcap) file now includes the session ID in the filename. We now store the file in **destination-IP-address.destination-port.protocol.session-id.pcap** format in the **/var/log/pcap** location. (Previously, the packet capture file was saved in **destination-IP-address. destination-port.protocol.pcap** format.)

  In addition, we've enhanced packet capture of unknown application functionality to capture unknown Server Name Indication (SNI) details.

[See Packet Capture of Unknown Application Traffic Overview.]

- **Application signature enhancements (NFX Series, SRX Series, and vSRX)**—Starting in Junos OS Release 21.1R1, we've introduced the following enhancements to application signatures:
  - Support for FTP data context propagation
  - Skipping of deep packet inspection (DPI) for the sessions offloaded by advanced policy-based routing (APBR) on application system cache (ASC) hit (when only APBR service is enabled).
  - Forceful installation of the application signature pack over the same version of signature pack.
  - Display (in the CLI command output) of the application signature pack release date.
  - Display (in the CLI command output) of the list of deprecated application signatures available in the installed signature pack.

  [See Predefined Application Signatures for Application Identification.]

Authentication and Access Control

- **Configure client information to connect to the JIMS server (cSRX, SRX300, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting with Junos OS Release 21.1R1, you can configure which specifc interface, source IP address or routing instance SRX should use for connecting to a JIMS server.
  [See Configuring the Connection to an SRX Series Device.]

Chassis

- **Layer 2 channel error alarm (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 21.1R1, you can configure a threshold limit for the L2 channel error count and set an alarm when the error count crosses the threshold using two new configuration statements. Use l2-channel-errorthreshold to set a threshold limit for the L2 channel error count, and use l2-channel-errors to set an alarm when the error count crosses the threshold at the [set chassis alarm] hierarchy level. This configuration generates a SNMP trap of the L2 channel error whenever the alarm is raised.

  [See l2-channel-error-threshold (chassis), l2-channel-errors (chassis), and show system alarms.]

Chassis Cluster

- **Support for NAT functionalities on multinode high availability (SRX5400, SRX5600, and SRX5800 with SPC3 card)**—Starting in Junos OS Release 21.1R1, we support the following NAT functionalities on HA nodes in multinode high availability:
  - IPv6 NAT (source NAT, destination NAT, and static NAT)

- NAT64 persistent NAT
- Logical and tenant systems NAT (source NAT, destination NAT, and static NAT)
- Port block allocation (PBA) and NAT logs.

[See Multinode High Availability, and NAT for User Logical Systems.]

- **Enabling and disabling control link (SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 21.1R1, you can enable or disable control links, using the following commands, to control the status of the cluster nodes and minimize failovers.
  - **delete chassis cluster control-interface node 0 disable**
  - **delete chassis cluster control-interface node 1 disable**
  - **set chassis cluster control-interface node0 disable**
  - **set chassis cluster control-interface node1 disable**

  [See chassis and fabric-options].

Ethernet Switching and Bridging

- **LLDP on routed and reth interfaces (SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 21.1R1, you can enable LLDP on all physical interfaces, including routed and redundant Ethernet (reth) interfaces. LLDP is a link-layer protocol used by network devices to advertise capabilities, identity, and other information to a LAN.

  [See LLDP Overview.]

EVPN

- **EVPN-VXLAN tunnel inspection (SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)—**Starting in Junos OS Release 21.1R1, we've introduced the following enhancements to the VXLAN support for SRX Series devices:
  - Support for SRX5000 line of devices in addition to the SRX4000 line and vSRX
  - Enhancements to tunnel inspection for VXLAN-encapsulated traffic by applying Layer 4 or Layer 7 security services to the tunnel traffic. The supported services are:
    - Application identification
    - IDP
    - Juniper Advanced Threat Prevention (ATP Cloud)

- Unified threat management (UTM)

Layer 7 security services provide application-level security and protect users from security threats through VXLAN tunnel.

[See Configuring Tunnel Traffic Inspection.]

- **Security policy enhancement for EVPN-VXLAN tunnel inspection (SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting in Junos OS Release 21.1R1, we've enhanced EVPN-VXLAN tunnel inspection by adding zone-level policy control for the inner traffic. When you create a policy that applies to the inner session created by VXLAN inner header, you can define the following parameters as match conditions for the inner traffic:

  - Source zone
  - Destination zone
  - URL category
  - Dynamic applications

  Additional matching criteria in the security policy provide granular control and extensibility to manage traffic.

  [See Configuring Tunnel Traffic Inspection.]

Flow-Based and Packet-Based Processing

- **Support for PowerMode IPsec (PMI) solution (SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800 with SPC3 cards, vSRX, and vSRX3.0) and GRE acceleration solution (SRX Series and NFX Series)**—Starting in Junos OS Release 21.1R1, we support the PMI and GRE acceleration solutions to improve the software-defined WAN (SD-WAN) performance.

  Table 18: Solutions and Details

Solution	How to Enable?
PMI	Include the `power-mode-ipsec` and `gre-performance-acceleration` statements at the `[edit security flow]` hierarchy level. **NOTE**: PMI supports both IPsec and GRE. In this case, traffic flows through the PMI data path.

Table 18: Solutions and Details *(Continued)*

Solution	How to Enable?
GRE acceleration	Include the `gre-performance-acceleration` statement at the `[edit security flow]` hierarchy level. **NOTE**: By default, `gre-performance-acceleration` is turned off. In this case, traffic flows through the GRE acceleration data path.

[See gre-performance-acceleration (Security Flow), flow (Security Flow), and show security flow status.]

- **Enhanced monitoring and troubleshooting of the flow session (SRX Series)**—Starting in Junos OS Release 21.1R1, we've introduced additional filters to the show security flow session operational command. The additional filters allow you to generate specified outputs in a list so that you can easily monitor the flow session. We've also introduced the show security flow session pretty and show security flow session plugins operational commands to view detailed information about the flow session.

  You can also trace the packet-drop information without committing the configuration using the monitor security packet-drop operational command. This command output is displayed on the screen until you press Ctrl+c or until the security device collects the requested number of packet drops. The command includes various filters to generate the output fields per your requirement.

  [See show security flow session, show security flow session pretty, show security flow session plugins, and monitor security packet-drop.]

- **Packet-based ECMP support for Express Path (SRX5400, SRX5600, and SRX5800)**—In earlier releases, Express Path supported only session-based ECMP traffic. Starting in Junos OS Release 21.1R1, Express Path also supports packet-based ECMP traffic from different network processors of the SRX Series device. In the packet-based ECMP mode, the SPU creates multiple network processor sessions on multiple network processors at a time. This feature is enabled by default.

  [See Express Path.]

High Availability

- **Distributed mode support for fast BFD failure detection (SRX1500, SRX4100, SRX4200, and SRX4600)**—Starting in Junos OS Release 21.1R1, we support distributed mode for BFD. This mode provides a faster BFD failure detection time of 3 x 300 ms. You enable distributed mode by configuring the BFD failure detection time to a value less than 500 ms. We support this feature for a standalone SRX Series device. It is not supported for chassis clusters.

> **NOTE**: SRX1500 devices run in dedicated mode if you've configured `set chassis dedicated-ukern-cpu`, regardless of the BFD failure detection time. You can enable distributed mode on SRX1500 devices only when dedicated mode is not enabled.

[See detection-time (BFD Liveness Detection) and Understanding Distributed BFD.]

Interfaces

- **Native VLAN ID configuration on the reth interface (SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, and SRX4600)** —Configuring the native VLAN ID on the redundant Ethernet (reth) interface enables the logical interface whose VLAN ID matches the native VLAN ID that is configured for that interface to accept untagged packets as well as tagged packets. Using the same logical interface with the native VLAN ID enabled ensures that any packet going out of that interface does not have a tag attached. Packets can be outbound control packets or transit data packets.

  [See native-vlan-id.]

Intrusion Detection and Prevention

- **Support for Perl-compatible regular expression (PCRE) version 8.40 (SRX Series and NFX Series)**—Starting in Junos OS Release 21.1R1, we’ve upgraded the codebase of intrusion detection and prevention (IDP) from PCRE version 5.40 to PCRE version 8.40. As PCRE version 8.40 supports new regex constructs, this upgrade enhances the capability of Junos OS IDP attack signatures to match regular expressions. With this upgrade, we've also addressed security vulnerabilities in the Junos OS PCRE codebase.

  [See pattern-pcre (Security IDP).]

- **Support for Snort IPS signatures (SRX Series and NFX Series)**—Starting in Junos OS Release 21.1R1, Juniper Networks IDP supports Snort IPS signatures. IDP secures your network by using signatures that help to detect attacks. Snort is an open-source intrusion prevention system (IPS). You can convert the Snort IPS rules into Juniper IDP custom attack signatures using the Juniper Integration of Snort Tool (JIST). These rules help detect malicious attacks.

  - JIST is included in Junos OS by default. The tool supports Snort version 2 and version 3 rules.
  - JIST converts the Snort rules with snort-ids into equivalent custom attack signatures on Junos OS with respective snort-ids as the custom attack names.
  - When you run the `request` command with Snort IPS rules, JIST generates `set` commands equivalent to the Snort IPS rules. Use the `request security idp jist-conversion` command to generate the `set` commands as CLI output. To load the `set` commands, use the `load set terminal` statement or copy

and paste the commands in the configuration mode, and then commit. You can then configure the existing IDP policy with the converted custom attack signatures.

- All the Snort IPS rule files that didn't get converted are written to **/tmp/jist-failed.rules**. The error log files generated during the conversion are written to **/tmp/jist-error.log**.
- To view the jist-package version, use the `show security idp jist-package-version` command.

[See Understanding Snort IPS Signatures, request security idp jist-conversion , and show security idp jist-package-version .]

Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud)

- **Server Message Block (SMB) protocol support for Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud) file inspection (SRX Series)**—Starting in Junos OS Release 21.1R1, SRX Series devices support the SMB protocol in advanced anti-malware (AAMW) file inspection. Use the `set services advanced-anti-malware policy` *`policy-name`* `smb` command to configure file inspection for the SMB protocol.

  [See advanced-anti-malware policy and show services advanced-anti-malware statistics.]

- **Support for configuring DNS sinkhole (SRX5000 line of devices)**—Starting in Junos OS Release 21.1R1, we support DNS sinkhole feature on the SRX5000 line of devices in addition to its existing support on SRX4000 line of devices and vSRX. You can configure DNS filtering to identify DNS requests for disallowed domains. You can either:

  - Block access to the domain by sending a DNS response that contains the IP address or fully qualified domain name (FQDN) of a sinkhole server. This ensures that when the client attempts to send traffic to the disallowed domain, the traffic instead goes to the sinkhole server.
  - Log the DNS request and reject access.

  [See dns-filtering.]

- **Support for username feed type in adaptive threat profiling (SRX Series devices and vSRX)**—Starting in Junos OS Release 21.1R1, you can add the user source identity (username) as a feed type in adaptive threat profiling. Use the `add-source-identity-to-feed` *`user-identity`* and `add-destination-identity-to-feed` *`user-identity`* commands at the `[edit security policies from-zone` *`zone-name`* `to-zone` *`zone-name`* `policy` *`policy-name`* `then [permit|deny|reject] application-services]` hierarchy level to configure the username feed type.

  [See security-intelligence (security policies), show services security-intelligence sec-profiling-feed status and show services security-intelligence category.]

- **Enhancements to alerts, alarms, and fallback options (SRX Series)**—Starting in Junos OS Release 21.1R1, we've enhanced the following alerts, alarms, and fallback options for failure conditions when you enroll SRX Series devices with Juniper ATP Cloud.

  - Add new SNMP traps for the following:
    - Advanced-anti-malware (AAMW)—`jnxJsAAMWChannelUp` and `jnxJsAAMWChannelDown`.
    - Encrypted traffic insights—`jnxJsSMSChannelUp` and `jnxJsSMSChannelDown`
    - Security intelligence (SecIntel)—`jnxJsSecIntelChannelUp` and `jnxJsSecIntelChannelDown`
  - Raise new alarms for AAMW, encrypted traffic insights, and SecIntel.
  - Add new fallback options for action control in case of failure conditions. Configure the fallback options at the `[edit services advanced-anti-malware policy` *`policy-name`* `]` hierarchy level.

  [See advanced-anti-malware policy.]

- **Support for Juniper ATP Cloud services in VXLAN tunnel inspection (SRX4000 line of devices, SRX5000 line of devices, and vSRX)**—Starting in Junos OS Release 21.1R1, the listed SRX Series devices and vSRX support Juniper ATP Cloud services such as AAMW and SecIntel in VXLAN tunnel traffic inspection. These services inspect the VXLAN traffic only if there is a security policy configured to perform the inspection. When you configure VXLAN tunnel inspection policies on an SRX Series device, the device scans the VXLAN tunnel traffic through AAMW and SecIntel services.

  [See tunnel-inspection and show security flow session.]

- **Policy-based threat profiling (SRX Series devices and vSRX)**—Starting in Junos OS Release 21.1R1, you can add the user source identity (username) to a security policy to generate security feeds.

  Juniper ATP Cloud service consolidates the generated feeds from SRX Series device and shares the duplicated results back with that security device. The security device uses the feeds to perform actions against the designated traffic. You can enable the security device to use the feeds by configuring security policies with the feeds as matching criteria. When traffic matches policy conditions, the device applies policy actions.

  [See Threat Profiling Support in Security Policy.]

Network Management and Monitoring

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C,**

QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

[See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

- **HMAC-SHA-2 authentication protocol support for users of SNMPv3 USM (MX Series and SRX Series)**—Starting in Junos OS Release 21.1R1, you can configure HMAC-SHA-2 authentication protocols for users of the SNMPv3 user-based security model (USM) with the following new CLI configuration statements:

  - authentication-sha224
  - authentication-sha256
  - authentication-sha384
  - authentication-sha512

  We've introduced these statements for local-engine users at `[edit snmp v3 usm local-engine user username]` and for remote-engine users at `[set snmp v3 usm remote-engine engine-id user user-name]`.

  [See authentication-sha224, authentication-sha256, authentication-sha348, and authentication-sha512.]

- **Log profiles and templates for customized logging (cSRX, SRX300, SRX320, SRX340, SRX345, SRX380, SRX550HM, SRX1500, SRX4100, SRX4200, SRX4600, SRX5800, and vSRX)**—Starting in Junos OS Release 21.1R1, you can configure log profiles and log templates for a policy. Use the configuration statement `profile` to select a log profile for a policy at the `[edit security log profile]` hierarchy level, and use the configuration statement `template` to select a predefined log template for a policy at the `[edit security log profile profile-name template]` hierarchy level. From this release, you can track the application tracking logs using the `set security application-tracking log-session-create`, `set security application tracking log-session-close`, `set security application-tracking session-update-interval`, `set security application-tracking no-volume-updates`, and `set services application-identification no-application statistics` commands. Unified threat management (UTM) features also support the log profiles and templates for customized logging.

  [See profile (security), application-tracking, application-identification, and show security log profile.]

Securing GTP and SCTP Traffic

- **Support for messages and message lists for aggregate rate limiting (SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, and SRX5800 with SPC2 and SPC3 cards, and vSRX)**—Starting in Junos OS Release 21.1R1, we support messages and message lists for aggregate rate limiting. To configure a message list, include the `message-list` *`msg-list-name`* `message` *`msg-number`* statement at the `[edit security gtp]` hierarchy level. To configure the default messages, use the `rate-limit default message {v0 | v1 | v2}` *`msg-list-name`* statement at the `[edit security gtp]` hierarchy level. Use the `show security gtp message-list` to display message-list profiles. Use the `show security gtp rate-limit default` to display default rate-limit messages.

  [See message-list, rate-limit (Aggregated rate limit), show security gtp message-list, and show security gtp rate-limit default.]

Services Applications

- **Support for RFC 2544-based benchmarking tests (SRX300, SRX320, SRX340, SRX345, SRX380, and SRX550HM)**—Starting in Junos OS Release 21.1R1, we support only the Layer 3 reflector function for these tests, with the following limitations:

  - `family inet` option; no other families are supported
  - IPv4 source and destination addresses for the tests

  RFC 2544 tests measure and demonstrate the service-level agreement (SLA) parameters before service activation. You can use the tests to measure throughput, latency, frame loss rate, and the number of back-to-back frames. [See Understanding Using Probes for Real-Time Performance Monitoring on M, T, PTX, MX, and SRX Series Routers .]

Software Installation and Upgrade

- **request system software status command (MX480, MX960, MX2010, MX2020, SRX1500, SRX4100, SRX4400, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, you can use the CLI command `request system software status` to view the status of the software package installation or uninstallation on the local Routing Engine.

Unified Threat Management (UTM)

- **Source address configuration for UTM services (SRX Series)**—Starting in Junos OS Release 21.1R1, you can configure the `source-address` option at the following hierarchy levels for the Enhanced Web Filtering (EWF) cloud service, websense redirect policy service, and antivirus and antispam scan services. Configuring `source-address` for these Unified Threat Management (UTM) services enhances the network disaster recovery.

- `[edit security utm default-configuration web-filtering juniper-enhanced server]`
- `[edit security utm default-configuration web-filtering websense-redirect server]`
- `[edit security utm feature-profile web-filtering websense-redirect profile profile name server]`
- `[edit security utm default-configuration anti-virus sophos-engine server]`

Antivirus and antispam services share the same `source-address` configuration under the Sophos engine server.

[See source-address.]

VPNs

- **Enhancements to increase traffic selector flexibility (SRX5400, SRX5600, and SRX5800)**—Starting in Junos OS Release 21.1R1, you can do the following to add flexibility to your traffic selectors in different deployment scenarios:
  - Configure the routing metric for a traffic selector.
  - Define the source port range, destination port range, and protocol for a traffic selector.
  - Define multiple terms within a traffic selector, instead of creating multiple traffic selectors (or child security associations or SAs) for a VPN. Each term comprises the local and remote IP prefixes, the source and destination port ranges, and the protocol identifier. You can use these parameters in a single IPsec SA negotiation. In earlier Junos OS releases, you configure each traffic selector with one set of local and remote IP prefixes to be used in an IPsec SA negotiation with a peer.

  This feature is supported only if the `junos-ike` package is installed in your device.

  We recommend you configure the same metric value if you define multiple traffic selectors under the same `[edit security ipsec vpn` *`vpn_name`*`]` hierarchy level with same value for *remote-ip ip-address/ netmask*. If you configure different metric values, then the metric value of the st0 route installed will be same as the traffic selector that is negotiated or installed first.

  [See traffic-selector and show security ipsec security-associations detail.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for SRX Series.

What's Changed in Release 21.1R3

IN THIS SECTION

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

  [See request system scripts refresh-from and cert-file (Scripts).]

Platform and Infrastructure

- **No support for PKI operational mode commands on the Junos Limited version (MX Series routers, PTX Series routers, and SRX Series devices)**—We do not support `request` , `show` , and `clear` PKI-related operational commands on the limited encryption Junos image ("Junos Limited"). If you try to execute PKI operational commands on a limited encryption Junos image, then an appropriate error message is displayed. The `pkid` process does not run on Junos Limited version image. Hence, the limited version does not support any PKI-related operation.

J-Web

- Changes to the Dashboard and Monitor pages (SRX Series): To improve the J-Web UI loading speed:
  - On the Dashboard page, we've removed the on-box reports related widgets.
  - On the Monitor > Maps and Charts > Traffic Map page, we've changed the default duration from "Last 1 hour" to Last "5 minutes."

What's Changed in Release 21.1R2

IN THIS SECTION

Network Management and Monitoring

- **Change in OID ifHighSpeed**–Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

  [See SNMP MIBs and Traps Supported by Junos OS.]

- **New output field added in show pfe statistics traffic command (SRX380)**—Starting in Junos OS Release, you?ll see Unicast EAPOL in the output of the show pfe statistics traffic command.

  [See show pfe statistics traffic.]

What's Changed in Release 21.1R1

IN THIS SECTION

Flow-Based and Packet-Based Processing

- **Self-generated IKE packets choose outgoing interface matching source IP address (SRX Series)**—A self-generated IKE packet always selects the ECMP outgoing interface that matches the source IP address. Note that we don't support filter-based forwarding for self-generated traffic with rerouting.

General Routing

- **Change in show security firewall-authentication jims operational command (SRX4600)**—Starting in Junos OS Release 21.1R1, the `show security firewall-authentication jims (statistics | display)` operational command includes the `display` option.

  [See show security firewall-authentication jims statistics.]

- **New output field added in `show pfe statistics traffic` command (SRX380)**—Starting in Junos OS Release 21.1R1, you'll see `Unicast EAPOL` in the output of the `show pfe statistics traffic` command.

  [See show pfe statistics traffic.]

- **Default MKA transmit interval (SRX380)**—On SRX380 devices, the default MACsec Key Agreement (MKA) transmit interval is 2000 milliseconds. If you deploy an SRX380 device with another security peer device with a MACsec secure link, you must change the MKA transmit interval on the peer device to 2000 milliseconds to match the new default MKA transmit interval of the SRX380 device.

  [See transmit-interval (MACsec).]

Intrusion Detection and Prevention

- **Intelligent offload state (SRX Series)**—We've introduced a new field in the `show security idp status` command to see the status of the IDP Intelligent offload.

  [See show security idp status.]

Junos XML API and Scripting

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX commit scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX commit scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **The `jcs:invoke()` function supports suppression of root login and logout events in system log files for SLAX event scripts (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The `jcs:invoke()` extension function supports the `no-login-logout` parameter in SLAX event scripts. If you include the parameter, the function does not generate and log `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages when the script logs in as root to execute the specified remote procedure call (RPC). If you omit the parameter, the function behaves as in earlier releases in which the root `UI_LOGIN_EVENT` and `UI_LOGOUT_EVENT` messages are included in system log files.

  [See invoke() Function (SLAX and XSLT).]

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Network Management and Monitoring

- **Support for specifying the YANG modules to advertise in the NETCONF capabilities and supported schema list (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—You can configure devices to emit third-party, standard, and Junos OS native YANG modules in the capabilities exchange of a NETCONF session by configuring the appropriate statements at the `[edit system services netconf hello-message yang-module-capabilities]` hierarchy level. In addition, you can specify

the YANG schemas that the NETCONF server should include in its list of supported schemas by configuring the appropriate statements at the `[edit system services netconf netconf-monitoring netconf-state-schemas]` hierarchy level.

[See hello-message and netconf-monitoring.]

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

User Interface and Configuration

- **Verbose format option to export JSON configuration data (ACX Series, EX Series, MX Series, PTX Series, QFX Series, and SRX Series)**—The Junos OS CLI exposes the `verbose` statement at the `[edit system export-format json]` hierarchy level. We changed the default format to export configuration data in JavaScript Object Notation (JSON) from `verbose` to `ietf` starting in Junos OS Release 16.1R1. You can explicitly specify the default export format for JSON configuration data by configuring the appropriate statement at the `[edit system export-format json]` hierarchy level. Although the `verbose` statement is exposed in the Junos OS CLI as of the current release, you can configure this statement starting in Junos OS Release 16.1R1.

  [See export-format.]

VPNs

- **Support for trace options log levels (SRX5400, SRX5600, and SRX5800)**—You can configure the log levels using the `level (all | error | info | notice | verbose | warning)` statement at the `edit security ike traceoptions` hierarchy level for troubleshooting the IKE issues.

  [See traceoptions].

- **View the traffic selector type for an IPsec tunnel (SRX Series and MX Series)**—You can run the `show security ipsec security-associations detail` command to display the traffic selector type for a VPN. The `show security ipsec security-associations detail` command displays `proxy-id` or `traffic-selector` as a value for the `TS Type` output field based on your configuration.

  [See show security ipsec security-associations.]

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for SRX Series devices.

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

Authentication and Access Control

- Unified access control drops packet at uac-plugin receives interest check event when jbuf is null. When flow get uac return value, the flow process would drop packet and log again. Normally when flow process did the log, it will check the flag on jbuf. The flag would flow this drop had been logged and then flow won't log again. Sometimes, the jbuf was not there hence no flag can be set so two logs were seen. PR1555850

Flow-Based and Packet-Based Processing

- For accelerated flows such as Express Path, the packet or byte counters in the session close log and show session output take into account only the values that accumulated while traversing the NP. PR1546430

General Routing

- In SRX380, MACsec show security macsec statistics command, when encryption-offset is enabled, the encrypted bytes and encrypted packets will include both encrypted and protected bytes. PR1534840
- Due to enhancements in AppID starting Junos OS Release 21.1R1, database files are not compatible with earlier releases. Hence, this issue is expected to be seen during downgrade from Junos OS Release 21.1R1 to earlier releases. PR1554490

VPNs

- In SPC2 and SPC3 mixed-mode HA deployments, tunnel per second (TPS) is getting affected while dead peer detection (DPD) is being served on existing tunnels. This limitation is due to a large chunk of CPU being occupied by infrastructure (gencfg) used by IKED to synchronize its DPD state to the backup nodes. PR1473482

Open Issues

IN THIS SECTION

Learn about open issues in this release 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.

Flow-Based and Packet-Based Processing

- Use an antireplay window size of 512 for IPv4 or IPv6 in fat-tunnel. The ESP sequence check might otherwise report out-of-order packets if the fat-tunnel parallel encryption is within 384 packets (12 cores * 32 packets in one batch). Hence, there are no out-of-order packets with 512 antireplay window size. PR1470637

General Routing

- The PKI CMPv2 (RFC 4210) client certificate enrolment does not properly work on SRX Series devices when using root-CA. PR1549954
- On all SRX Series devices with Sky Advanced Threat Prevention (Sky ATP) used, when putting upper-case letters into the Realm name field during the Sky ATP CLI enrollment, it will output "Provided password for * is incorrect", which is incorrect and misleading. PR1550387
- Kernel might stop, with VM core files generated, and the system might reboot continuously after five child interfaces are added to the reth interface on one node. This might cause service impact. PR1551297
- When the device is downgraded to a release earlier than Junos OS Release 21.1 and then upgraded again to Junos OS Release 21.1, the appiddb tables might not get populated properly and have 0 entries. For such cases, after upgrading, uninstall and reinstall signature package. PR1567199
- PKID core might occur during cert signature validation . This core is not very frequent and occurs due to memory corruption . PR1573892
- With ssl-proxy configured along with web-proxy, the client session might not closed on the device even though proxy session ends gracefully. PR1580526
- Web-proxy: Getting UNKNOWN instead of HTTP-PROXY for application and UNKNOWN instead of GOOGLE-GEN in RT-FLOW close messages These messages can be seen in the RT-flow close log and these are due to JDPI not engaged for the session. This may affect the app identification for the web-proxy session traffic. PR1588139
- On SRX345, icmp checksum error and packet drops are observed while doing rapid ping on vdsl interface with MTU 1514. PR1591230

- There is a behaviour change in application track logs. By default, logs are disabled. PR1591966
- In Junos OS releases 20.3R3, 20.4R3 and 21.1R2, sometimes on reboot schedule report are not getting generated. PR1594377
- For Junos OS releases 20.3R3, 20.4R3, 21.1R2, 21.2R1, phone home ZTP is failing on SRX Series devices as phone home client is unable to connect to Phone Home Server or Redirect Server. PR1598462
- When static routes are added with gr interface names, there could be replication issues with mpls nexthops causing backup to core. PR1601996

Interfaces and Chassis

- Traffic drop might be seen on irb interface on SRX1500 for network control forwarding class when verifying dscp classification based on single and multiple code-points. PR1611623

Intrusion Detection and Prevention (IDP)

- On SRX Series devices, it is unable to use latest signature pack due to IDP DB failing to update. PR1594283
- IDPD will not core when wrong package in given for offline download and it will do two level of validation.
  - Look for mandatory file in offline downloaded Package.
  - Secpack having manifest files which contains the list of files to be expected in package.

  So the fix is based on above file if package is missing any file from manifest file list then package will be considered as bad package. PR1623857

Platform and Infrastructure

- On SRX Series devices with Bidirectional Forwarding Detection (BFD) enabled for multiple protocols (such as OSPF, ISIS, BGP, PIM), the ppmd process might crash after an upgrade. PR1335526
- If authentication (tacplus-server, radius-server) is configured on a device, it may fail to open files in a rare case, which may cause the process mgd to stop. PR1600615

Routing Policy and Firewall Filters

- If a huge number of policies are configured on SRX Series devices and some policies are changed, the traffic that matches the changed policies might be dropped. PR1454907
- When SSL Proxy's global-config is set with with enable-proxy-on-default-fw-policy-match, the traffic is hitting pre-id policy instead of default policy for Yahoo traffic. PR1542790

Routing Protocols

- Commit error seen while adding static route for a link-local IPv6 destination address range. PR1599273

VPNs

- When multiple traffic selectors are configured on a particular VPN, the iked process checks for a maximum of 1 DPD probe that is sent to the peer for the configured DPD interval. The DPD probe is sent to the peer if traffic flows over even one of the tunnels for the given VPN object. PR1366585
- In the output of the show security ipsec inactive-tunnels command, Tunnel Down Reason is not displayed as this functionality is not supported in Junos OS Release 18.2R2 and later. PR1383329
- On SRX5400, SRX5600, and SRX5800 devices with an SPC3 card, a new behavior has been introduced that differs from the behavior on the older SPC2 card. The SRX Series device with AutoVPN configuration can now accept multiple IPsec tunnels from a peer device (with the same source IP address and port number) using different IKE IDs. PR1407356
- On SRX5400, SRX5600, and SRX5800 devices, during in-service software upgrade (ISSU), the IPsec tunnels flap, causing a disruption of traffic. The IPsec tunnels recover automatically after the ISSU process is completed. PR1416334
- In some scenario(e.g configuring firewall filter) sometimes srx5K might show obsolete IPsec SA and NHTB entry even when the peer tear down the tunnel. PR1432925
- An IPsec policy must not have both ESP and AH proposals. The configuration will commit, but the IPsec traffic will not work. Do not configure an IPsec policy with proposals using both ESP and AH protocols. PR1552701
- Do not configure two traffic selectors for the same peer under the same IPsec VPN with the same values. PR1554533

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases 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: 21.1R3

IN THIS SECTION

Authentication and Access Control

- UAC authentication might not work post system reboot. PR1585158

Flow-Based and Packet-Based Processing

- Security traffic log display service-name="None" for some application. PR1619321

General Routing

- SSL-FP Logging for non SNI session. PR1442391
- Wi-Fi mPIM on SRX Series devices is reaching out to NTP and DNS servers. PR1569680
- When using log templates (introduced in Junos OS release21.1R1) with Unified Policies, logs were not generated in a predictable manner. A new construct has been added that allows you to define a default log profile (set security log profile *name* default-profile) that can be used to improve this behaviour when multiple log profiles are defined. PR1570105
- Missing snmp operation state method for on SRX5800 or MX960 devices. PR1570433
- Changes in SNMP traps configuration and data exported for TWAMP. PR1573169
- On SRX Series devices, error message tcp_timer_keep:Local(0x81100001:60753) Foreign(0x8f100001:33010) is seen in messages log every 80 seconds. PR1580667
- Traffic is dropped to or through VRRP virtual IP on SRX380 devices. PR1581554
- The srxpfe process might stop on SRX1500 devices. PR1582989
- Secure Web proxy continue sending DNS query for unresolved DNS entry even after the entry was removed. PR1585542
- On SRX Series devices, significant performance improvements for JDPI's micro-application identification were included. PR1585683
- IP packets might be dropped on SRX Series devices. PR1588627
- The jsqlsyncd process files generation might cause device to stop after upgrade. PR1589108
- The REST API does not work for SRX380 devices. PR1590810
- In Junos OS releases 20.1R3 and 20.3R3, the issue (empty feed-name) starts with the hit returned from cache which points to the node with the parameter of feed-ID (2) inconsistent with the feeds-update (when it's 1). As a result the incorrect feed-ID points to the empty entry in the array of the feed-names. PR1591236

- J-Web Deny log nested-application="UNKNOWN" instead of specific application. PR1593560
- When combining log profiles and unified policies RT_FLOW_SESSION_DENY logs were not being generated corrected. PR1594587
- When JDPI inspection-limits are reached, under certain circumstances, classification details were not propagated to interested Layer-7 Services, such as IDP. PR1595310
- Node1 fpc0 (SPM) goes down after ISSU and RG0 failover. PR1595462
- Jflow V9 application-id record: Network based application recognition value for IPv4 application-id are not as expected. PR1595787
- Delay might be observed between Services Processing Card (SPC) failing and failover to other node. PR1596118
- The flowd might core dump if application-services security policy is configured. PR1597111
- The srxpfe process might stop and generate a core file post "targeted-broadcast forward-only" interface-config commit. PR1597863
- The flowd process might generate core files, if the AppQOS module receiving two packets of a session. PR1597875
- The flowd process might stop in AppQoE scenarios. PR1599191
- The httpd-gk core might be observed when IPsec VPN is configured. PR1599398
- Traffic might be dropped at NAT gateway if EIM is enabled. PR1601890
- In the best path log message the switch reason is being shown as nh change instead of sla violated. PR1602571
- The flowd process might stop if the DNS-inspection feature is enabled by configuring SMS policy. PR1604773
- Memory leak at the useridd process might be observed when integrated user firewall is configured. PR1605933
- When the tap mode is enabled, the packet on ge-0/0/0 is dropped on RX side. PR1606293
- DNS proxy functionality might not work on VRRP interfaces. PR1607867
- Enabling security-metadata-streaming-policy might cause Packet Forwarding Engine to stop.PR1610260
- Interface might not come up when 10G port is connected to 1G SFP. PR1613475

Interfaces and Chassis

- IPv4 or IPv6 address from the config on the interface may not be applied when the interface is moved from tenants to interface stanza in the configuration. PR1605250

Intrusion Detection and Prevention (IDP)

- Custom attack IDP policies might fail to compile. PR1598867
- IDP policy compilation is not happening when a commit check is issued prior to a commit. PR1599954
- The srxpfe might stop while the IDP security package contains a new detector. PR1601380
- Optimizations made to IDP that help improve its performance and behaviour under load. PR1601926
- High Routing Enginer CPU usage occurs when routing-instance is configured under security idp security-package hierarchy level. PR1614013

J-Web

- J-Web may not display customer defined application services if one new policy is created. PR1599434
- J-web application might stop and generates httpd process core files. PR1602228
- Radius users might not be able to view or modify configuration through J-Web. PR1603993
- On all SRX Series devices, some widgets in J-Web might not load properly for logical systems users. PR1604929
- The J-Web error: "your session has expired. click ok to re-login" when using root user. PR1611448

Network Address Translation (NAT)

- Incorrect IPv6 UDP checksum inserted after translation of packet from IPv4 to IPv6. PR1596952

Platform and Infrastructure

- SPC3 might not come up after the system reboot. PR1555904
- Junos OS: Upon receipt of specific sequences of genuine packets destined to the device the kernel will crash and restart (vmcore) (CVE-2021-0283, CVE-2021-0284). PR1595649

- On SRX Series devices, the accounting and auditd process on secondary node does not work. PR1620564

Routing Policy and Firewall Filters

- High CPU usage might be seen on some SRX Series devices. PR1579425

Routing Protocols

- Short multicast packets drop using PIM when multicast traffic received at a non-RPT or SPT interface. PR1579452

User Interface and Configuration

- After image upgrade device might fail to come up due to certain configurations. PR1585479

VPNs

- The iked process might restart and generate core during session state activation or deactivation. PR1573102
- Memory leaks on the iked process on SRX5000 line of devices with SRX5K-SPC3 installed. PR1586324
- The IPSec tunnel might not come up if configured with configuration payload in a certain scenario. PR1593408
- The kmd process might crash when VPN peer initiates using source-port other than 500. PR1596103
- Tail drops might occur on SRX Series devices if shaping-rate is configured on st-interface. PR1604039

Resolved Issues: 21.1R2

IN THIS SECTION

EVPN

- The mustd process generates core files during upgrading or while committing a configuration. PR1577548

Flow-Based and Packet-Based Processing

- The flowd or srxpfe process might crash when clearing the TCP proxy session. PR1573842
- On SRX Series devices, the filter from-zone has been added to the utility monitor security packet-drop. PR1574060

General Routing

- The flowd might generate core files frequently on SRX340 device. PR1463689
- The kmd process might crash when the interface flaps. PR1544800
- SRX1500 reports fans running at over speed. PR1546132
- Application identity unknown packet capture utility does not function on SRX Series devices when enhanced-services mode is enabled. PR1558812
- The PIC in SRX5K-SPC3 or MX-SPC3 card might get stuck in offline status after flowd crash occurs on it. PR1560305
- Fabric probe packets might be processed incorrectly when power-mode-ipsec is enabled. PR1564117
- Wi-Fi mPIM on SRX Series devices is reaching out to NTP and DNS servers. PR1569680
- Missing SNMP operation state method for on SRX5800 and MX960 devices. PR1570433

- MACsec not using network-control queue. PR1571977
- Traffic going through the VRRP interface might be dropped when VRRP enabled IRB interface goes down. PR1572920
- In certain conditions where on SRX , the timer values are updated for an existing fast BFD session , it may cause a fast BFD session deletion on the PFE. This will result in BFD session remaining DOWN or PFE core occasionally. PR1578946
- The ipfd process might crash with a coredump when SecProfiling thread feeds are fetched from Policy Enforcer(PE) PR1582454
- The srxpfe process might stop on SRX1500 device. PR1582989
- The 1G interfaces might not come up after device reboot. PR1585698
- The l2ald process might crash on changing the routing instance. PR1586516
- On SRX Series devices, the protocol-version command which controls TLS-versions (1.1, 1.2, 1.3, etc) within SSL proxy has been unhidden. PR1587149
- On SRX Series devices, the unknown packet capture functionality will no longer record SSL. The UNKNOWN flows by default. This behaviour can be changed by enabling the command set services application-identification packet-capture ssl-unknown. Without configuration the ssl-unknown command, the SRX Series device will only capture flows marked as UNKNOWN or INCONCLUSIVE. PR1587875
- Pass through traffic might fail post reboot when Secure Web Proxy is configured. PR1589957

Interfaces and Chassis

- When the interface configured under tenant instance and an interface configured under default routing instance have identical local address. While trying to delete tenant, the interface under it is getting added to default routing instance, which is causing commit check error for identical local address under same routing instance. As a workaround, commit full instead of commit when deleting tenant instance.PR1581877

Intrusion Detection and Prevention (IDP)

- Adding signature in packet drop reason and sending to record packet drops module. PR1574603
- The IDP policy process might become unresponsive and fail to compile the IDP policy after an IDP automatic update. PR1577684

J-Web

- To improve performance in Monitoring > Network > Interfaces page, Admin Status is removed, Services and Protocols data merged into one Host inbound traffic. PR1574895
- The zone info disappears when functional zone is configured. PR1594366
- A custom application name contains any is listed under pre-defined applications. PR1597221

Network Address Translation (NAT)

- Incorrect IPv6 UDP checksum inserted after translation of packet from IPv4 to IPv6. PR1596952

Network Management and Monitoring

- SSH connection might become unresponsive and logs show kern.maxfiles limit exceeded by uid messages. PR1567634

Platform and Infrastructure

- The show chassis errors command is not supported on SRX5000 line of devices with RE3 and SCB3 installed anymore. PR1560562
- The show chassis ethernet-switch errors command unexpectedly shows error counters for port 14 on the SRX5800 device. PR1563978
- On SRX5000 line of devices, the alarm Power Budget:Insufficient Power may be raised incorrectly when the second SCB does not contain an Routing Engine. PR1568183

Routing Policy and Firewall Filters

- Traffic loss might be seen when a big number of applications or addresses is referenced by one policy. PR1576038

VPNs

- The pkid process generates core files while auto-enrollment of local certificates. PR1564300
- When there are multiple IPsec SAs, backup SA starts IPsec rekey. PR1565132
- The iked process might crash by operational commands on the SRX5000 line of devices with SRX5000-SPC3 card installed. PR1566649
- The kmd process might crash when VPN peer initiates using source-port other than 500. PR1596103

Resolved Issues: 21.1R1

Chassis Clustering

- Disabled node on SRX cluster sent out ARP request packets. PR1548173
- SPU process stop might be seen under a GPRS tunneling protocol scenario. PR1559802

Flow-Based and Packet-Based Processing

- When no logical system or tenant system flow trace is configured and no root-override is configured, the latest behavior is to not log any flow trace for that logical system or tenant system, instead of dumping all to root flow trace as before. PR1530904
- THR capacity update on SRX Series devices. PR1538058
- The rst-invalidate-session command does not work if configured together with the no-sequence-check command. PR1541954
- Application fragmented traffic might get dropped on SRX Series devices. PR1543044
- Instability with RGs on cluster. PR1550637
- Adjust the default route change timeout value. PR1553621
- The usp_max_tcplib_connection is not expected on SRX1500, SRX4100, and SRX4200 devices. PR1563881

General Routing

- On the SRX1500 device, the traffic rate shown in the CLI command is not accurate. PR1527511
- The MAC table is null in Layer 2 mode after one pass-through session is created successfully. PR1528286
- The firewall filter SA and DA tags are not in the log messages as expected in port details. PR1539338
- Packet drop might be seen when a packet with destination port 0 is received on the SRX380 device. PR1540414
- Tail drops might occur on SRX Series devices if shaping-rate is configured on lt- interface. PR1542931
- The nsd process might stop when DNS-based allowlisting is configured under SSL proxy. PR1542942

- The Wi-Fi Mini-Physical Interface Module (Mini-PIM) does not support pure g mode with 2.4-GHz radio. PR1543824
- The output of the show services application-identification group detail command incorrectly included Micro-Applications (Micro-Apps) in the output of every group. PR1544727
- On SRX4100 and SRX4200 devices, if PEM0 is removed, the output of jnxOperatingDescr.2 might be incomplete. PR1547053
- Advanced anti-malware file or e-mail statistics does not get incremented with the latest PB version. PR1547094
- Continuous "LCC: ch_cluster_lcc_set_context:564: failed to lock chassis_vmx mutex 11" chassisd logs generated. PR1547953
- Lcmd log "gw_cb_presence:136: PEM(slot = 0): error detecting presence (fruid = 15, drv_id = 30, status = -11)" generated every second on the SRX4100 and SRX4200. PR1550249
- On SRX1500, SRX-SFP-1GE-T (Part#740-013111) for a copper cable might be corrupted after reboot. PR1552820
- The volume displayed in traffic map are redefined. PR1553066
- The speed mismatch error is seen while trying to commit reth0 with gigether-options. PR1553888
- An IPFD core file might be generated when using Adaptive Threat Profiling. PR1554556
- On an SRX550M device, the dumpdisklabel command fails with message "ERROR: Unknown platform srx550m." PR1557311
- AppID's Unknown Packet Capture utility does not function on SRX Series devices when enhanced-services mode is enabled. PR1558812
- The show security log report top session-close group-by application order-by risk top-number 8 where-application-risk high xml encapsulation structure changed and caused script fail. PR1559013
- The show security log report top idp group-by threat-severity order-by count top-number 5 where-attack command display will change the idp reporting to match the threat-severity in idp log.. PR1560027
- High CPU usage on pkid process might be seen when the device is unable to connect to a particular CRL URL. PR1560374
- The DNS commands may not be executed and also any new configuration may not take effect on connecting the SRX Series device to Juniper ATP Cloud. PR1561169
- There is an idpd core file at ../../../../../../../src/junos/secure/usr.sbin/idp-confd/idpd_lsys.c:771. PR1561298

- When multiple IRB interfaces belong to the same VRRP group ID, if one of IRB intefaces goes down, it causes disruption in traffic going through another IRB interface. PR1572920

Interfaces and Chassis

- When SRX Series devices receive proxy ARP requests on VRRP interfaces, the devices send ARP replies with the underlying interface MAC address. PR1526851
- Backup Routing Engine or backup node may be stuck in bad status with an improper backup-router configuration. PR1530935

Intrusion Detection and Prevention (IDP)

- The greater than or less than symbols are allowed for age-of-attack filter of dynamic attack group configuration. The age-of-attack field in signatures will be changed to CVE dates from activation dates.

  PR1397599
- The flowd or srxpfe process might generate core files during the idpd process commit on SRX Series devices. PR1521682
- IDP now supports the ability to create dynamic-attack-groups based on attack-prefix wildcards. For example, you can include all of the Metasploit-based scans by applying this filter to a dynamic-attack-group: set attack-prefix values SCAN:METASPLOIT:*. PR1537195
- SOF support for partial packet plugins on traditional or unified policy. PR1542497
- Need syslog to indicate signature download completion. PR1543571
- IDP policy load might fail post image upgrade for Junos OS Release 15.1X49 releases. PR1546542
- The idpd process crashes and generates a core file. PR1547610

J-Web

- Sometimes, when you edit the local gateway in the remote access VPN workflow under VPN>IPsec VPN, J-web might not display one or more drop-down values. PR1521788
- J-Web browser tab title to include product model name and hostname. PR1523760
- J-Web GUI does not allow you to save the rules with more than 2500 cumulative shared objects. PR1540047
- After commit pending changes message is shown, the contents of other messages, landing page, or pop-ups will not be visible completely. PR1554024

Layer 2 Ethernet Services

- The RG1 interface failover occurs when RG0 failover is triggered. PR1366825

Platform and Infrastructure

- Syslog reporting PFE_FLOWD_SELFPING_PACKET_LOSS: Traffic impact: Selfping packets loss/err: 300 within 600 second error messages in node 0 and node 1 control panel. PR1522130
- The commit might not fail as expected when reth interface is deleted. PR1538273

Routing Policy and Firewall Filters

- Traffic might be dropped unexpectedly when the url-category match condition is used on a security policy. PR1546120
- Global policies working with multi-zones cause high Packet Forwarding Engine CPU utilization. PR1549366
- Policy configured with the route-active-on condition may work incorrectly for local routes. PR1549592
- NSD process stops when the secprofiling feed name is 64 bytes. PR1549676
- The junos-defaults construct within a unified-policies application match criteria now restricts the ports and protocols of a flow on a per-dynamic-application basis. PR1551984
- Unified policies in global zone contexts do not work when from-zone or to-zone is defined. PR1558009
- On the SRX5000 line of devices, the secondary node might get stuck in performing ColdSync after a reboot or upgrade, or if ISSU is performed. PR1558382
- The traffic may be dropped if you insert one global policy above others on SRX Series devices. PR1558827

Subscriber Access Management

- Incorrect counter type (counter instead of gauge) specified for some values in MIB jnxUserAAAMib. PR1533900

Unified Threat Management (UTM)

- Stream buffer memory leak might happen when UTM is configured under unified policies. PR1557278

- UTM license expiry event lost may cause the device can't quit advance service mode and maximum-sessions decreased by half. PR1563874

User Interface and Configuration

- The outbound-ssh routing-instance is shown as unsupported. PR1558808

VPNs

- The output of show security ipsec security-associations command might display empty space instead of keyword null for encryption algorithm. PR1507270
- On all SRX Series devices using IPsec with NAT traversal, MTU size for the external interface might be changed after IPsec SA is reestablished. PR1530684
- After IPsec tunnel using policy-based VPN is overwritten by another VPN client, traffic using this IPsec tunnel will be dropped. PR1546537
- Traffic going through a policy-based IPsec tunnel might be dropped after RG0 failover. PR1550232
- The iked process may crash with L3HA setup. PR1559121
- The iked process might crash by operational commands on the SRX5000 line of devices with SRX5000-SPC3 card installed. PR1566649

Documentation Updates

This section lists the errata and changes in Junos OS Release 21.1R3 for the SRX Series documentation.

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains the upgrade and downgrade support policy for Junos OS for SRX Series devices. 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.

For information about ISSU, see the Chassis Cluster User Guide for Security Devices.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 19: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Junos OS Release Notes for vMX

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for vMX. 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for vMX.

What's New in 21.1R3

There are no new features for vMX in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features for vMX in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for vMX.

Network Management and Monitoring

- **Ephemeral configuration database support for load update operations (EX9200, MX5, MX10, MX80, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, and vMX)**—Starting in Junos OS Release 21.1R1, NETCONF and Junos XML protocol client applications can configure the ephemeral configuration database using a load update operation. To perform a load update operation, set the `<load-configuration>` `action` attribute to `update`.

  [See <load-configuration>.]

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, operational command RPCs, including the `<get-configuration>` RPC, support the `format="json-minified"` and `format="xml-minified"` attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

[See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

Software Installation and Upgrade

- **request system software status command (MX480, MX960, MX2010, MX2020, SRX1500, SRX4100, SRX4400, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, you can use the CLI command `request system software status` to view the status of the software package installation or uninstallation on the local Routing Engine.

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for vMX.

What's Changed in Release 21.1R3

IN THIS SECTION

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the

server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

[See request system scripts refresh-from and cert-file (Scripts).]

Platform and Infrastructure

- **Enhancement to the request system license add terminal command (PTX10001-36MR and vMX** — When you run the `request system license add terminal` command. You can now view following additional fields for `information: JUNOS564022985: Ignoring unknown feature .`

  [See Managing vMX Licenses.]

What's Changed in Release 21.1R2

Platform and Infrasturcture

- We are discontinuing support for vMX on Microsoft Azure starting in Junos OS Release 18.1R1.

What's Changed in Release 21.1R1

IN THIS SECTION

Junos XML API and Scripting

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Network Management and Monitoring

- **Change in license bandwidth command on vMX virtual routers**

  —Starting in Junos OS, to use the available license bandwidth, explicitly set the license bandwidth use the set chassis license bandwidth <In Mbps> command.

  [See Configuring Licenses on vMX Virtual Routers.]

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

Known Limitations

There are no known limitations for vMX in Junos OS Release 21.1R3.

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

Open Issues

IN THIS SECTION

Learn about open issues in this release for vMX.

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

Platform and Infrastructure

- The port speed shows as 100G even though chassis config is set for 40G. This is just a cosmetic display issue. PR1623237
- On vMX, the blockpointer in the ktree is getting corrupted leading to core-file generation. There is no function impact such as fpc restart or system down and the issue is not seen in hardware setups. PR1525594

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for vMX.

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

Resolved Issues: 21.1R3

IN THIS SECTION

Interfaces and Chassis

- Interface hold-time up does not work on vMX and MX150. PR1604554

Resolved Issues: 21.1R2

There are no resolved issues for vMX in Junos OS Release 21.1R2.

Resolved Issues: 21.1R1

General Routing

- Multiple vmxt core files might be generated on vMX platforms. PR1534641
- The riot forwarding process pause might be seen on vMX platforms configured with an IRB interface. PR1544856
- Observed ping failure on vMX while verifying SCU accounting. PR1569047

Upgrade Instructions

You cannot upgrade Junos OS for the vMX router from earlier releases using the `request system software add` command.

You must deploy a new vMX instance using the downloaded software package.

Remember to prepare for upgrades with new license keys and/or deploying Agile License Manager.

Junos OS Release Notes for vRR

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for vRR. 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.

What's New

IN THIS SECTION

There are no new features for vRR in Junos OS Release 21.1R1.

What's New in 21.1R3

There are no new features for vRR in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features for vRR in Junos OS Release 21.1R2.

What's New in 21.1R1

Learn about new features or enhancements to existing features in this release for vRR.

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for vRR.

What's Changed in Release 21.1R3

There are no changes in behavior or syntax for vRR in Junos OS Release 21.1R3.

What's Changed in Release 21.1R2

There are no changes in behavior or syntax for vRR in Junos OS Release 21.1R2.

What's Changed in Release 21.1R1

There are no changes in behavior or syntax for vRR in Junos OS Release 21.1R1.

Known Limitations

There are no known limitations for vRR in Junos OS Release 21.1R3.

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

To learn more about common BGP or routing known limitations in Junos OS 21.1R3, see "Known Limitations" on page 147 for MX Series routers.

Open Issues

There are no known issues for vRR in Junos OS Release 21.1R3.

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

To learn more about common BGP or routing knowns issues in Junos OS 21.1R3, see "Open Issues" on page 152 for MX Series routers.

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for vRR.

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

Resolved Issues: 21.1R3

IN THIS SECTION

To learn more about common BGP or routing resolved issues in Junos OS 21.1R3, see "Resolved Issues: 21.1R3" on page 175 for MX Series routers.

General Routing

- Memory might be exhausted when both the BGP rib-sharding and the BGP ORR (Optimal Route Reflection) enabled. PR1613104
- The process rpd might crash in BGP rib-sharding scenario. PR1613723

Resolved Issues: 21.1R2

IN THIS SECTION

To learn more about common BGP or routing resolved issues in Junos OS 21.1R2, see "Resolved Issues: 21.1R2" on page 192 for MX Series routers.

Platform and Infrastructure

- JRR200: Option-60 (Vendor-Class-Identifier) is not sent during ZTP. PR1582038

Resolved Issues: 21.1R1

To learn more about common BGP or routing resolved issues in Junos OS 21.1R1, see "Resolved Issues: 21.1R1" on page 213 for MX Series routers.

Routing Protocols

- BGP flap and rpd crash might be observed. PR1545837
- 6PE prefixes may not be removed from the RIB upon reception of withdrawal from a BGP neighbor when RIB sharding is enabled. PR1556271

Junos OS Release Notes for vSRX

IN THIS SECTION

These release notes accompany Junos OS Release 21.1R3 for vSRX. 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.

What's New

IN THIS SECTION

Learn about new features introduced in the Junos OS main and maintenance releases for vSRX.

What's New in 21.1R3

There are no new features for vSRX in Junos OS Release 21.1R3.

What's New in 21.1R2

There are no new features for vSRX in Junos OS Release 21.1R2.

What's New in 21.1R1

IN THIS SECTION

Learn about new features or enhancements to existing features in this release for vSRX.

Authentication and Access Control

- **Configure client information to connect to the JIMS server (cSRX, SRX300, SRX340, SRX345, SRX1500, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, and vSRX)**—Starting with

Junos OS Release 21.1R1, you can configure which specifc interface, source IP address or routing instance SRX should use for connecting to a JIMS server.

[See Configuring the Connection to an SRX Series Device.]

- **LLDP support in Layer 3 mode (vSRX 3.0)**—Starting in Junos OS Release 21.1R1, vSRX 3.0 in Layer 3 mode supports Link Layer Discovery Protocol (LLDP) to learn and distribute device information on network links. The device information enables the vSRX 3.0 to identify a variety of devices quickly. This quick identification results in a LAN that interoperates smoothly and efficiently.

  [See Device Discovery Using LLDP and LLDP-MED on Switches and lldp.]

Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud)

- **Support for username feed type in adaptive threat profiling (SRX Series devices and vSRX)**—Starting in Junos OS Release 21.1R1, you can add the user source identity (username) as a feed type in adaptive threat profiling. Use the `add-source-identity-to-feed` *`user-identity`* and `add-destination-identity-to-feed` *`user-identity`* commands at the `[edit security policies from-zone` *`zone-name`* `to-zone` *`zone-name`* `policy` *`policy-name`* `then [permit|deny|reject] application-services]` hierarchy level to configure the username feed type.

  [See security-intelligence (security policies), show services security-intelligence sec-profiling-feed status and show services security-intelligence category.]

- **Support for Juniper ATP Cloud services in VXLAN tunnel inspection (SRX4000 line of devices, SRX5000 line of devices, and vSRX)**—Starting in Junos OS Release 21.1R1, the listed SRX Series devices and vSRX support Juniper ATP Cloud services such as AAMW and SecIntel in VXLAN tunnel traffic inspection. These services inspect the VXLAN traffic only if there is a security policy configured to perform the inspection. When you configure VXLAN tunnel inspection policies on an SRX Series device, the device scans the VXLAN tunnel traffic through AAMW and SecIntel services.

  [See tunnel-inspection and show security flow session.]

- **Policy-based threat profiling (SRX Series devices and vSRX)**—Starting in Junos OS Release 21.1R1, you can add the user source identity (username) to a security policy to generate security feeds.

  Juniper ATP Cloud service consolidates the generated feeds from SRX Series device and shares the duplicated results back with that security device. The security device uses the feeds to perform actions against the designated traffic. You can enable the security device to use the feeds by configuring security policies with the feeds as matching criteria. When traffic matches policy conditions, the device applies policy actions.

  [See Threat Profiling Support in Security Policy.]

Licensing

- **Juniper Agile Licensing (vSRX)**—
  Starting in Junos OS Release 21.1R1, we're moving toward supporting license-based software features. We now use Juniper Agile Licensing to support soft enforcement for virtual CPU (vCPU) usage on vSRX. With soft enforcement, you can use more vCPUs than the number of vCPU licenses you are entitled to use. However, if you do that, the device generates an alarm. You can see the list of alarms at System Log Explorer.

  Juniper Agile Licensing provides simplified and centralized license administration and deployment. You can use Juniper Agile Licensing to install and manage licenses for hardware and software features.

  "Table 20" on page 387 describes the licensing support for vSRX.

  Table 20: Licensed Features on vSRX

vSRX License Model	Use Case Examples or Solutions	Number of vCPUs Required	Feature List
Standard	Use for standard firewall and secure branch routers	2, 5, 9, 17, or 32 virtual CPUs (vCPUs)	Application Layer Gateways (ALGs), BGP, class of service (CoS), DHCP, diagnostics, firewall, GRE, IP tunneling, IPv4 and IPv6, J-Flow, management (J-Web, CLI, and NETCONF), MPLS, multicast, NAT, on-box logging, OSPF, screens, site-to-site VPN, static routing, and user firewall
Advanced	Advanced 1 Use for data center security	2, 5, 9, 17, or 32 vCPUs	Includes Standard features plus IPS and application security (application identification, application firewall, application quality of service, and application tracking)
	Advanced 2 Use for next-generation firewall with cloud-based antivirus	2, 5, 9, 17, or 32 vCPUs	Includes Standard and Advanced 1 features, Sophos antivirus, Web filtering, antispam, and content filtering

Table 20: Licensed Features on vSRX *(Continued)*

vSRX License Model	Use Case Examples or Solutions	Number of vCPUs Required	Feature List
	Advanced 3 Use for next-generation firewall with on-box antivirus	2, 5, 9, 17, or 32 vCPUs	Includes Standard and Advanced 1 features, Avira antivirus, Web filtering, antispam, and content filtering
Premium	Premium 1 Use for data center security and Juniper Advanced Threat Prevention Cloud (Juniper ATP Cloud)	2, 5, 9, 17, or 32 vCPUs	Includes Standard and Advanced 1 features, and Juniper ATP Cloud
	Premium 2 Use for next-generation firewall and Juniper ATP Cloud	2, 5, 9, 17, or 32 vCPUs	Includes Standard and Advanced 2 features, and Juniper ATP Cloud
	Premium 3 Use for next-generation firewall and Juniper ATP Cloud	2,5,9, 17, or 32 vCPUs	Includes Standard and Advanced 3 features, and Juniper ATP Cloud

[See Flex Software License for vSRX, Juniper Agile Licensing Guide, and Configuring Licenses in Junos OS.]

Network Management and Monitoring

- **Operational command RPCs support returning JSON and XML output in minified format in NETCONF sessions (ACX1000, ACX1100, ACX2100, ACX4000, ACX5048, ACX5096, ACX5448, EX2300, EX3400, EX4300, EX4400-24P, EX4400-24T, EX4400-48F, EX4400-48P, EX4400-48T, EX4600, EX4650, EX9200, MX104, MX204, MX240, MX480, MX960, MX2008, MX2010, MX2020, MX10003, MX10008, MX10016, PTX1000, PTX5000, PTX10001, PTX10002, PTX10008, PTX10016, QFX5100, QFX5110, QFX5120, QFX5200, QFX5210, QFX10002, QFX10002-60C, QFX10008, QFX10016, SRX550HM, SRX4100, SRX4200, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, operational command RPCs, including the

<get-configuration> RPC, support the format="json-minified" and format="xml-minified" attributes in NETCONF sessions to return JSON or XML output in minified format. Minified format removes any characters that are not required for computer processing—for example, unnecessary spaces, tabs, and newlines. Minified format decreases the size of the data, and as a result, can reduce transport costs as well as data delivery and processing times.

[See Specifying the Output Format for Operational Information Requests in a NETCONF Session.]

Software Installation and Upgrade

- **Phone-home client (vSRX)**—Starting in Junos OS Release 21.1R1, the phone-home client (PHC) is responsible for the initial bootup and configuration of the vSRX VM instance when the virtual machine (VM) instance is turned on. When the vSRX VM instance boots up with the factory-default configuration, the phone-home client connects to a redirect server, which then redirects to the phone-home serverr. The phone-home client downloads the initial configuration and the latest Junos OS image from the phone-home server. .The new image is installed first, and then the initial configuration is applied and committed on the vSRX VM instance.

  If the redirect server does not provide any phone-home server information, the phone-home client restarts the provisioning process and keeps connecting to the redirect server until provisioning is successful.

  [See Obtaining Configurations and Software Image Without User Intervention Using Phone-Home Client.]

- **request system software status command (MX480, MX960, MX2010, MX2020, SRX1500, SRX4100, SRX4400, SRX4600, SRX5400, SRX5600, SRX5800, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, you can use the CLI command request system software status to view the status of the software package installation or uninstallation on the local Routing Engine.

VPNs

- **Increased tunnel scaling (vSRX 3.0)**—Starting in Junos OS Release 21.1R1, vSRX 3.0 is supported by a new architecture similar to SRX5000 line of devices with SPC3 which increases the tunnel scale.

  vSRX 3.0 instances support the IPsec VPN features that are supported on the SRX5000 line of devices with SPC3 (SRX5K-SPC3).

  By default, when the vSRX 3.0 boots up, the legacy architecture is executed. To enable the new architecture, you must load and install a new package, junos-ike. The Junos OS releases includes this package, but its installation is optional. As an administrator, you must execute the request system software add optional://junos-ike.tgz command to load the junos-ike package.

  [See IPsec VPN Features and Configurations Not Supported on SRX5K-SPC3 and vSRX Instances.]

What's Changed

IN THIS SECTION

Learn about what changed in the Junos OS main and maintenance releases for vSRX.

What's Changed in Release 21.1R3

IN THIS SECTION

Junos XML API and Scripting

- **Refreshing scripts from an HTTPS server requires a certificate (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—When you refresh a local commit, event, op, SNMP, or Juniper Extension Toolkit (JET) script from an HTTPS server, you must specify the certificate (Root CA or self-signed) that the device uses to validate the server's certificate, thus ensuring that the server is authentic. In earlier releases, when you refresh scripts from an HTTPS server, the device does not perform certificate validation.

  When you refresh a script using the `request system scripts refresh-from` operational mode command, include the `cert-file` option and specify the certificate path. Before you refresh a script using the `set refresh` or `set refresh-from` configuration mode command, first configure the `cert-file` statement under the hierarchy level where you configure the script. The certificate must be in Privacy-Enhanced Mail (PEM) format.

  [See request system scripts refresh-from and cert-file (Scripts).]

What's Changed in Release 21.1R2

IN THIS SECTION

Network Management and Monitoring

- **Change in OID ifHighSpeed**–Now, the object identifier (OID) ifHighSpeed displays the negotiated speed once negotiation is completed. If the speed is not negotiated, ifHighSpeed displays the actual maximum speed of the interface. In earlier releases, ifHighSpeed always displayed the actual speed of the interface.

  [See SNMP MIBs and Traps Supported by Junos OS.]

What's Changed in Release 21.1R1

IN THIS SECTION

Junos XML API and Scripting

- **Python 2.7 deprecation (ACX Series, EX Series, MX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—Starting in Junos OS Release 21.1R1, devices running Junos OS no longer support Python 2.7. We've deprecated the corresponding `language python` statement at the `[edit system scripts]` hierarchy level. To execute Python scripts, configure the `language python3` statement at the `[edit system scripts]` hierarchy level to execute the scripts using Python 3.

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

Network Management and Monitoring

- **Support for disconnecting unresponsive NETCONF-over-SSH clients (ACX Series, EX Series, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, and vSRX)**—You can enable devices to automatically disconnect unresponsive NETCONF-over-SSH clients by configuring the `client-alive-interval` and `client-alive-count-max` statements at the `[edit system services netconf ssh]` hierarchy level. The `client-alive-interval` statement specifies the timeout interval in seconds, after which, if no data has been received from the client, the device requests a response. The `client-alive-count-max` statement specifies the threshold of missed client-alive responses that triggers the device to disconnect the client, thereby terminating the NETCONF session.

  [See ssh (NETCONF).]

Upgrade and Downgrade

- **vCPU Core-Based License Requirement**—Starting in Junos OS Release 21.1R1, vSRX 3.0 requires a vCPU core-based license when you perform a software upgrade. When the software is upgraded, a warning message is displayed in Syslog indicating the requirement of the new vCPU core-based license.

  [See KB37351.]

Known Limitations

IN THIS SECTION

Learn about known limitations in this release for vSRX.

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

General Routing

- For SaaS DBs among all available links a best path chosen. If the link has no violation, and is the preferred link and has the highest prioty among all live links, any further configuration change won't be recognized. The recommendation to the user is to configure all the preferences and priorities during configuration time so that all of it can be properly honored. PR1559662

Software License

- The changes to vSRX3.0 license infrastructure, which caused Jflow v5 functionality to break in Routing Engine. PR1549988

Open Issues

IN THIS SECTION

Learn about open issues in this release for vSRX.

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

General Routing

- The vSRX cluster on KVM using Mellanox may go into a state where FPC is online but PIC is in ready state, this might cause cold sync failure and instabilities. PR1467342
- Intermittent license check core observed during the device initialization. License process will restart and start providing the required support. There is no service impact. PR1545175

- Tag "RT_FLOW_SESSION_XXX" is missing in stream mode. PR1565153
- When the device is downgraded to a release earlier than Junos OS Release 21.1 and then upgraded again to Junos OS Release 21.1, the appiddb tables might not get populated properly and have 0 entries. For such cases, after upgrading, uninstall and reinstall signature package. PR1567199
- With ssl-proxy configured along with web-proxy, the client session might not closed on the device even though proxy session ends gracefully. PR1580526
- Getting UNKNOWN instead of HTTP-PROXY for application and UNKNOWN instead of GOOGLE-GEN in RT-FLOW close messages These messages can be seen in the RT-flow close log and these are due to JDPI not engaged for the session. This may affect the app identification for the web-proxy session traffic. PR1588139

Routing Policy and Firewall Filters

- when SSL Proxy's global-config is set with with enable-proxy-on-default-fw-policy-match, the traffic is hitting pre-id policy instead of default policy for Yahoo traffic. PR1542790

VPNs

- In certain cases, the PUSH ACK message from the group member to the group key server may be lost. The group member can still send rekey requests for the TEK SAs before the hard lifetime expiry. Only if the key server sends any new PUSH messages to the grooup members, those updates would not be received by the group member since the key server would have removed the member from registered members list. PR1608290

Resolved Issues

IN THIS SECTION

Learn which issues were resolved in the Junos OS main and maintenance releases for vSRX.

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

Resolved Issues: 21.1R3

IN THIS SECTION

Application Layer Gateways (ALGs)

- ALG traffic might be dropped. PR1598017

Authentication and Access Control

- UAC authentication might not work post system reboot. PR1585158

Flow-Based and Packet-Based Processing

- Multicast traffic drop may occur on TAP interface on SRX Series devices. PR1583214

General Routing

- When using log templates (introduced in Junos OS release 21.1R1) with Unified Policies, logs were not generated in a predictable manner. A new construct has been added that allows you to define a default log profile (set security log profile *name* default-profile) that can be used to improve this behaviour when multiple log profiles are defined. PR1570105
- vSRX unreachable over SSH after integration with KMS on AWS. PR1584415
- When combining log profiles and unified policies RT_FLOW_SESSION_DENY logs were not being generated corrected. PR1594587

- Jflow V9 application-id record: Network based application recognition value for IPv4 application-id are not as expected. PR1595787
- In the best path log message the switch reason is being shown as nh change instead of sla violated. PR1602571
- vSRX might stop forwarding traffic 60 days after Junos upgrade due to the trial license expiring. PR1609551
- For apps getting classified on first packet, the volume update syslog is not getting generated. PR1613516
- The interface speed is limited to 1G on vSRX 2.0 even the speed is set as more than 1G. PR1617397

Network Address Translation (NAT)

- The object "jnxJsNatSrcNumPortAvail" does not show the proper value. PR1611479

Resolved Issues: 21.1R2

IN THIS SECTION

Flow-Based and Packet-Based Processing

- The flowd or srxpfe process might crash when clearing the TCP proxy session. PR1573842

General Routing

- The Jflow v5 functionality will not work correctly due to presence of new license infrastructure that is ported recently to vSRX3.0. PR1549988

- IKE configure mode payload is not pushing secondary DNS and secondary WINS attributes to Xauth module with IKEv1. Hence, the client is not getting assigned with secondary DNS and secondary WINS with IKEv1. PR1558831
- Delay in vSRX CLI prompt might be observed. PR1559741
- Fabric probe packets might be processed incorrectly when power-mode-ipsec (PMI) is enabled. PR1564117
- The rpd process generates core files at boot time of a device. PR1567043
- The srxpfe process might stop and generate a core file during the feed update process. PR1579631
- When a vSRX was performing DNS sinkholing, the sinkhole response packets that it would generate had incorrect checksums. This would cause the receiving client to drop the packet and not be directed to the vSRX's sinkhole. PR1582827

Intrusion Detection and Prevention (IDP)

- Global data SHM utilization increase quickly and FTP traffic might impacted. PR1585485
- Application identification related signatures might not get triggered. PR1588450

J-Web

- To improve performance in Monitoring > Network > Interfaces page, Admin Status is removed, Services and Protocols data merged into one Host inbound traffic. PR1574895

Platform and Infrastructure

- COS queue egress interface forwarding-class might not work as expected. PR1538286
- If committing source-address addr routing-instance and then delete source-address addr in private edit mode, commit fails with warning message. PR1582529

Routing Protocols

- Traffic might be lost during mirror data transmit from the primary ppmd or bfdd. PR1570228

VPNs

- When there are multiple IPsec SAs, backup SA starts IPsec rekey. PR1565132

Resolved Issues: 21.1R1

General Routing

- The control link might be broken when there is excessive traffic load on the control link in a vSRX cluster deployment. PR1524243
- The master-password configuration is rejected if master-encryption-password (MEK) is not set. PR1537251
- The srxpfe process might crash when the Application Identification Packet-Capture functionality is enabled. PR1538991
- Upgrading to Junos OS Release 20.4R1 or later releases with a large, preexisting security-log database might result in LLMD consuming large amounts of CPU. PR1548423
- Configuration integrity mismatch error in vSRX3.0 running on Azure with key-vault integrated. PR1551419
- The command set protocols l2-learning global-mode is removed on vSRX3.0. Use the show ethernet-switching global-information command. PR1554388
- High CPU usage on pkid process might be seen when the device is unable to connect to a particular CRL URL. PR1560374

Intrusion Detection and Prevention (IDP)

- The flowd or srxpfe process might generate core files during the idpd process commit on SRX Series devices. PR1521682
- On vSRX3.0 the attack-group-entries filters direction 0 limit 1 command is not showing expected values. PR1564761

J-Web

- The J-Web GUI does not allow you to save the rules with more than 2500 cumulative shared objects. PR1540047
- After commit pending changes message is shown, the contents of other messages, landing page, or pop-upsare not visible completely. PR1554024

Migration, Upgrade, and Downgrade Instructions

IN THIS SECTION

This section contains information about how to upgrade Junos OS for vSRX using the CLI. Upgrading or downgrading Junos OS can take several hours, depending on the size and configuration of the network.

You also can upgrade to Junos OS Release 21.1R3 for vSRX using J-Web (see J-Web) or the Junos Space Network Management Platform (see Junos Space).

Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Releases 17.4, 18.1, 18.2, 18.3,18.4, 19.1, 19.2 and 19.4 is supported.

The following limitations apply:

- Direct upgrade of vSRX from Junos OS 15.1X49 Releases to Junos OS Release 19.3 and higher is not supported. For upgrade between other combinations of Junos OS Releases in vSRX and vSRX 3.0, the general Junos OS upgrade policy applies.
- The file system mounted on /var usage must be below 14% of capacity.

  Check this using the following command:

```
show system storage | match " /var$" /dev/vtbd1s1f
 2.7G        82M       2.4G        3%  /var
```

  Using the `request system storage cleanup` command might help reach that percentage.
- The Junos OS upgrade image must be placed in the directory /var/host-mnt/var/tmp/. Use the `request system software add /var/host-mnt/var/tmp/<upgrade_image>`
- We recommend that you deploy a new vSRX virtual machine (VM) instead of performing a Junos OS upgrade. That also gives you the option to move from vSRX to the newer and more recommended vSRX 3.0.
- Ensure to back up valuable items such as configurations, license-keys, certificates, and other files that you would like to keep.

NOTE: For ESXi deployments, the firmware upgrade from Junos OS Release 15.1X49-Dxx to Junos OS releases 17.x, 18.x, or 19.x is not recommended if there are more than three network adapters on the 15.1X49-Dxx vSRX instance. If there are more than three network adapters and you want to upgrade, then we recommend that you either delete all the additional network adapters and add the network adapters after the upgrade or deploy a new vSRX instance on the targeted OS version.

Upgrading Software Packages

To upgrade the software using the CLI:

1. Download the **Junos OS Release 21.1R3 for vSRX .tgz** file from the Juniper Networks website. Note the size of the software image.
2. Verify that you have enough free disk space on the vSRX instance to upload the new software image.

```
root@vsrx> show system storage
       Filesystem              Size       Used      Avail  Capacity   Mounted on
       /dev/vtbd0s1a           694M       433M       206M       68%  /
       devfs                   1.0K       1.0K         0B      100%  /dev
       /dev/md0                1.3G       1.3G         0B      100%  /junos
       /cf                     694M       433M       206M       68%  /junos/cf
       devfs                   1.0K       1.0K         0B      100%  /junos/dev/
       procfs                  4.0K       4.0K         0B      100%  /proc
       /dev/vtbd1s1e           302M        22K       278M        0%  /config
       /dev/vtbd1s1f           2.7G        69M       2.4G        3%  /var
       /dev/vtbd3s2             91M       782K        91M        1%  /var/host
       /dev/md1                302M       1.9M       276M        1%  /mfs
       /var/jail               2.7G        69M       2.4G        3%  /jail/var
       /var/jails/rest-api       2.7G        69M       2.4G      3%  /web-api/var
       /var/log                2.7G        69M       2.4G        3%  /jail/var/log
       devfs                   1.0K       1.0K         0B      100%  /jail/dev
       192.168.1.1:/var/tmp/corefiles       4.5G       125M       4.1G    3%  /var/crash/
corefiles
       192.168.1.1:/var/volatile       1.9G       4.0K       1.9G    0%  /var/log/host
       192.168.1.1:/var/log       4.5G       125M       4.1G     3%  /var/log/hostlogs
       192.168.1.1:/var/traffic-log       4.5G       125M       4.1G    3%  /var/traffic-log
       192.168.1.1:/var/local       4.5G       125M       4.1G    3%  /var/db/host
```

```
        192.168.1.1:/var/db/aamwd       4.5G       125M       4.1G     3%  /var/db/aamwd
        192.168.1.1:/var/db/secinteld       4.5G       125M       4.1G     3%  /var/db/secinteld
```

3. Optionally, free up more disk space, if needed, to upload the image.

```
root@vsrx> request system storage cleanup
       List of files to delete:
       Size Date       Name
       11B Sep 25 14:15 /var/jail/tmp/alarmd.ts
       259.7K Sep 25 14:11 /var/log/hostlogs/vjunos0.log.1.gz
       494B Sep 25 14:15 /var/log/interactive-commands.0.gz
       21.1K Sep 25 14:15 /var/log/messages.0.gz
       27B Sep 25 14:15 /var/log/wtmp.0.gz
       27B Sep 25 14:14 /var/log/wtmp.1.gz
       3027B Sep 25 14:13 /var/tmp/BSD.var.dist
       0B Sep 25 14:14 /var/tmp/LOCK_FILE
       666B Sep 25 14:14 /var/tmp/appidd_trace_debug
       0B Sep 25 14:14 /var/tmp/eedebug_bin_file
       34B Sep 25 14:14 /var/tmp/gksdchk.log
       46B Sep 25 14:14 /var/tmp/kmdchk.log
       57B Sep 25 14:14 /var/tmp/krt_rpf_filter.txt
       42B Sep 25 14:13 /var/tmp/pfe_debug_commands
       0B Sep 25 14:14 /var/tmp/pkg_cleanup.log.err
       30B Sep 25 14:14 /var/tmp/policy_status
       0B Sep 25 14:14 /var/tmp/rtsdb/if-rtsdb
       Delete these files ? [yes,no] (no) yes
<
output omitted>
```

NOTE: If this command does not free up enough disk space, see [SRX] Common and safe files to remove in order to increase available system storage for details on safe files you can manually remove from vSRX to free up disk space.

4. Use FTP, SCP, or a similar utility to upload the Junos OS Release 21.1R3 for vSRX .tgz file to **/var/crash/corefiles/** on the local file system of your vSRX VM. For example:

```
root@vsrx> file copy ftp://username:prompt@ftp.hostname.net/pathname/
junos-vsrx-x86-64-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE.tgz /var/crash/corefiles/
```

5. From operational mode, install the software upgrade package.

```
root@vsrx> request system software add /var/crash/corefiles/junos-vsrx-
x86-64-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE.tgz no-copy no-validate reboot
Verified junos-vsrx-x86-64-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE signed by
PackageDevelopmentEc_2017 method ECDSA256+SHA256
THIS IS A SIGNED PACKAGE
WARNING:     This package will load JUNOS 21.1 software.
WARNING:     It will save JUNOS configuration files, and SSH keys
WARNING:     (if configured), but erase all other files and information
WARNING:     stored on this machine.  It will attempt to preserve dumps
WARNING:     and log files, but this can not be guaranteed.  This is the
WARNING:     pre-installation stage and all the software is loaded when
WARNING:     you reboot the system.
Saving the config files ...
Pushing Junos image package to the host...
Installing /var/tmp/install-media-srx-mr-vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE.tgz
Extracting the package ...
total 975372
-rw-r--r-- 1 30426 950 710337073 Oct 19 17:31 junos-srx-mr-
vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-app.tgz
-rw-r--r-- 1 30426 950 288433266 Oct 19 17:31 junos-srx-mr-
vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-linux.tgz
Setting up Junos host applications for installation ...
============================================
Host OS upgrade is FORCED
Current Host OS version: 3.0.4
New Host OS version: 3.0.4
Min host OS version required for applications: 0.2.4
============================================
Installing Host OS ...
upgrade_platform: --------------------
upgrade_platform: Parameters passed:
upgrade_platform: silent=0
upgrade_platform: package=/var/tmp/junos-srx-mr-vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-
linux.tgz
upgrade_platform: clean install=0
upgrade_platform: clean upgrade=0
upgrade_platform: Need reboot after staging=0
upgrade_platform: --------------------
upgrade_platform:
upgrade_platform: Checking input /var/tmp/junos-srx-mr-
```

```
vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-linux.tgz ...
upgrade_platform: Input package /var/tmp/junos-srx-mr-
vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-linux.tgz is valid.
upgrade_platform: Backing up boot assets..
cp: omitting directory '.'
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK
initrd.cpio.gz: OK
upgrade_platform: Checksum verified and OK...
/boot
upgrade_platform: Backup completed
upgrade_platform: Staging the upgrade package - /var/tmp/junos-srx-mr-
vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-linux.tgz..
./
./bzImage-intel-x86-64.bin
./initramfs.cpio.gz
./upgrade_platform
./HOST_COMPAT_VERSION
./version.txt
./initrd.cpio.gz
./linux.checksum
./host-version
bzImage-intel-x86-64.bin: OK
initramfs.cpio.gz: OK
version.txt: OK
upgrade_platform: Checksum verified and OK...
upgrade_platform: Staging of /var/tmp/junos-srx-mr-
vsrx-21.1-2021-07-07.0_RELEASE_21.1_THROTTLE-linux.tgz completed
upgrade_platform: System need *REBOOT* to complete the upgrade
upgrade_platform: Run upgrade_platform with option -r | --rollback to rollback the upgrade
Host OS upgrade staged. Reboot the system to complete installation!
WARNING:     A REBOOT IS REQUIRED TO LOAD THIS SOFTWARE CORRECTLY. Use the
WARNING:     'request system reboot' command when software installation is
WARNING:     complete. To abort the installation, do not reboot your system,
WARNING:     instead use the 'request system software rollback'
WARNING:     command as soon as this operation completes.
NOTICE: 'pending' set will be activated at next reboot...
Rebooting. Please wait ...
shutdown: [pid 13050]
Shutdown NOW!
*** FINAL System shutdown message from root@ ***
System going down IMMEDIATELY
```

```
Shutdown NOW!
System shutdown time has arrived\x07\x07
```

If no errors occur, Junos OS reboots automatically to complete the upgrade process. You have successfully upgraded to Junos OS Release 21.1R3 for vSRX.

NOTE: Starting in Junos OS Release 17.4R1, upon completion of the vSRX image upgrade, the original image is removed by default as part of the upgrade process.

6. Log in and use the `show version` command to verify the upgrade.

```
--- JUNOS 21.1-2021-07-07.0_RELEASE_21.1_THROTTLE Kernel 64-bit
JNPR-11.0-20171012.170745_fbsd-
At least one package installed on this device has limited support.
Run 'file show /etc/notices/unsupported.txt' for details.
root@:~ # cli
root> show version
Model: vsrx
Junos: 21.1-2021-07-07.0_RELEASE_21.1_THROTTLE
JUNOS OS Kernel 64-bit  [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs [20171012.170745_fbsd-builder_stable_11]
JUNOS OS runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS OS time zone information [20171012.170745_fbsd-builder_stable_11]
JUNOS OS libs compat32 [20171012.170745_fbsd-builder_stable_11]
JUNOS OS 32-bit compatibility [20171012.170745_fbsd-builder_stable_11]
JUNOS py extensions [20171017.110007_ssd-builder_release_174_throttle]
JUNOS py base [20171017.110007_ssd-builder_release_174_throttle]
JUNOS OS vmguest [20171012.170745_fbsd-builder_stable_11]
JUNOS OS crypto [20171012.170745_fbsd-builder_stable_11]
JUNOS network stack and utilities [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Web Management Platform Package [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs compat32 [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx runtime [20171017.110007_ssd-builder_release_174_throttle]
JUNOS common platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx platform support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS mtx network modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS modules [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srxtvp modules [20171017.110007_ssd-builder_release_174_throttle]
```

```
JUNOS srxtvp libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx libs [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx Data Plane Crypto Support [20171017.110007_ssd-builder_release_174_throttle]
JUNOS daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS srx daemons [20171017.110007_ssd-builder_release_174_throttle]
JUNOS Online Documentation [20171017.110007_ssd-builder_release_174_throttle]
JUNOS jail runtime [20171012.170745_fbsd-builder_stable_11]
JUNOS FIPS mode utilities [20171017.110007_ssd-builder_release_174_throttle]
```

Validating the OVA Image

If you have downloaded a vSRX .ova image and need to validate it, see Validating the vSRX .ova File for VMware.

Note that only .ova (VMware platform) vSRX images can be validated. The .qcow2 vSRX images for use with KVM cannot be validated the same way. File checksums for all software images are, however, available on the download page.

Upgrade and Downgrade Support Policy for Junos OS Releases and Extended End-Of-Life Releases

We have two types of releases, standard EOL and EEOL:

- Standard End of Life (EOL) releases have engineering support for twenty four months after the first general availability date and customer support for an additional six more months.
- Extended End of Life (EEOL) releases have engineering support for thirty six months after the first general availability date and customer support for an additional six more months.

For both standard EOL and EEOL releases, you can upgrade to the next three subsequent releases or downgrade to the previous three releases. For example, you can upgrade from 20.4 to the next three releases – 21.1, 21.2 and 21.3 or downgrade to the previous three releases – 20.3, 20.2 and 20.1.

For EEOL releases only, you have an additional option - you can upgrade directly from one EEOL release to the next two subsequent EEOL releases, even if the target release is beyond the next three releases. Likewise, you can downgrade directly from one EEOL release to the previous two EEOL releases, even if the target release is beyond the previous three releases. For example, 20.4 is an EEOL release. Hence, you can upgrade from 20.4 to the next two EEOL releases – 21.2 and 21.4 or downgrade to the previous two EEOL releases – 20.2 and 19.4.

Table 21: EOL and EEOL Releases

Release Type	End of Engineering (EOE)	End of Support (EOS)	Upgrade/ Downgrade to subsequent 3 releases	Upgrade/ Downgrade to subsequent 2 EEOL releases
Standard End of Life (EOL)	24 months	End of Engineering + 6 months	Yes	No
Extended End of Life (EEOL)	36 months	End of Engineering + 6 months	Yes	Yes

For more information about standard EOL and EEOL releases, see https://www.juniper.net/support/eol/junos.html.

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

Licensing

In 2020, Juniper Networks introduced a new software licensing model. The Juniper Flex Program comprises a framework, a set of policies, and various tools that help unify and thereby simplify the multiple product-driven licensing and packaging approaches that Juniper Networks has developed 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 about the list of supported products, see Juniper Flex Program.

Finding More Information

- **Feature Explorer**—Juniper Networks Feature Explorer helps you to explore software feature information to find the right software release and product for your network.

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

- **PR Search Tool**—Keep track of the latest and additional information about Junos OS open defects and issues resolved.

  https://prsearch.juniper.net/InfoCenter/index?page=prsearch

- **Hardware Compatibility Tool**—Determine optical interfaces and transceivers supported across all platforms.

  https://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.

  https://pathfinder.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:

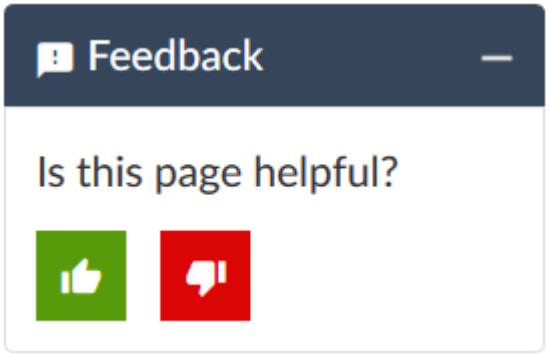

- Click the thumbs-up icon if the information on the page was helpful to you.
- Click the thumbs-down icon if the information on the page was not helpful to you or if you have suggestions for improvement, and use the pop-up form to provide feedback.

- 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

IN THIS SECTION

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

1 June 2023—Revision 8, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

10 February 2023—Revision 7, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

25 November 2022—Revision 6, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

24 March 2022—Revision 5, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

17 March 2022—Revision 4, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

3 March 2022—Revision 3, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

28 January 2022—Revision 2, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

27 December 2021—Revision 1, Junos OS Release 21.1R3– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

11 November 2021—Revision 4, Junos OS Release 21.1R2– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

7 October 2021—Revision 3, Junos OS Release 21.1R2– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

9 September 2021—Revision 2, Junos OS Release 21.1R2– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

6 August 2021—Revision 1, Junos OS Release 21.1R2– ACX Series, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

24 June 2021—Revision 8, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

10 June 2021—Revision 7, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

13 May 2021—Revision 6, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

29 April 2021—Revision 5, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

7 April 2021—Revision 4, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

31 March 2021—Revision 3, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

29 March 2021—Revision 2, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.

26 March 2021—Revision 1, Junos OS Release 21.1R1– ACX Series, cRPD, cSRX, EX Series, JRR Series, Juniper Secure Connect, Junos Fusion, MX Series, NFX Series, PTX Series, QFX Series, SRX Series, vMX, vRR, and vSRX.